AD-A043 246 ILLINOIS UNIV AT URBANA-CHAMPAIGN SAVOY AVIATION RES--ETC F/G 5/9
AN EVALUATION OF A COMPLEX COMPUTER-BASED FLIGHT PROCEDURES TRA--ETC(U)
JAN 77 S R TROLLIP F44620-76-C-0009

UNCLASSIFIED ARL-77-1/AFOSR-77-1 AFOSR-TR-77-1013 NL

1 OF 2
AD
A043 246

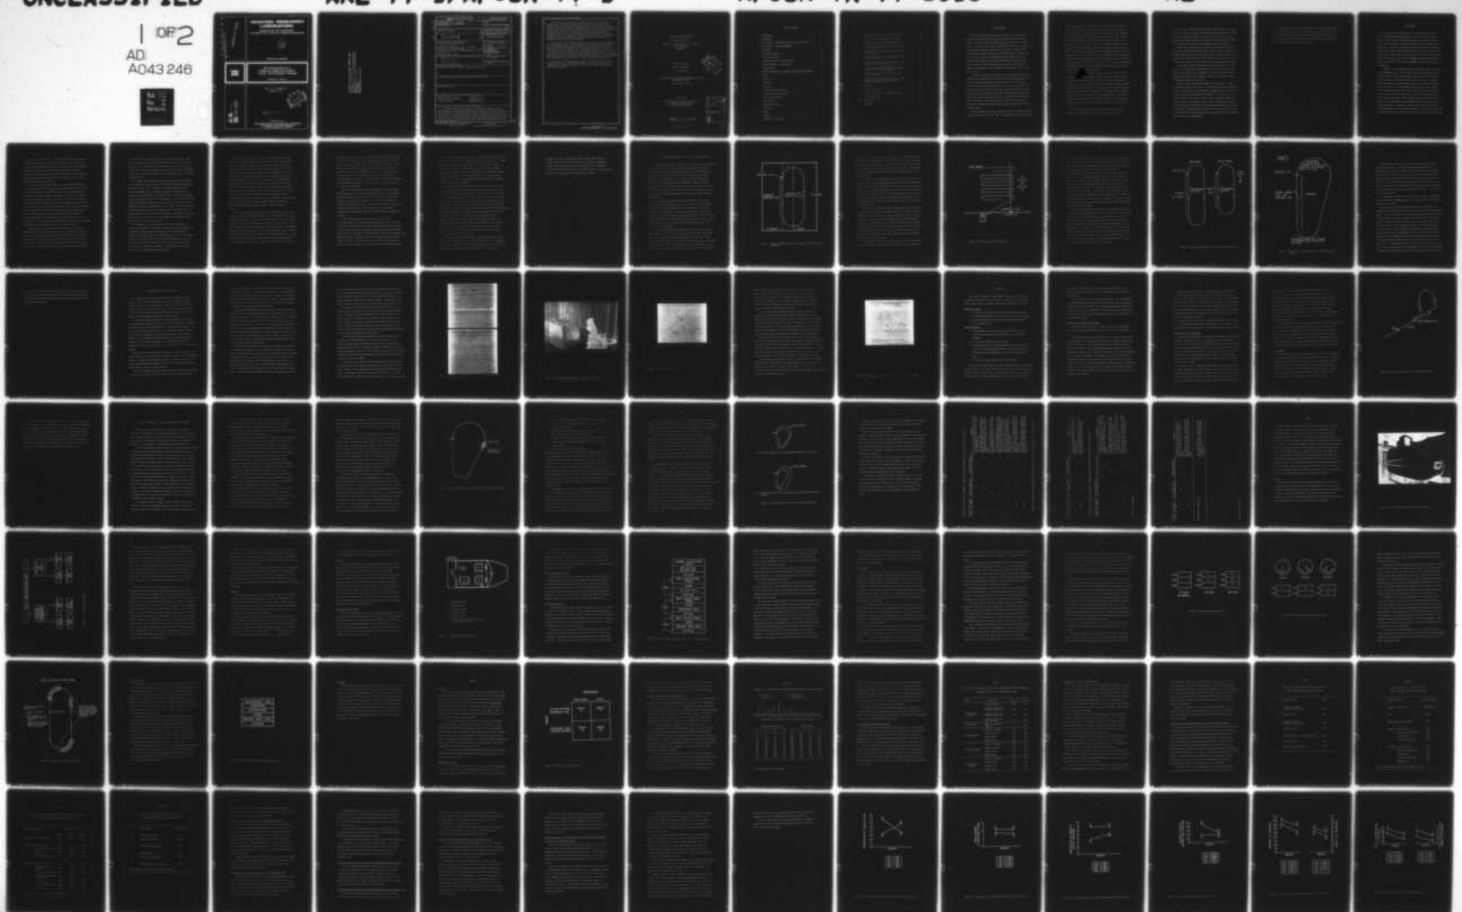

AD A043246

AFOSR-TR- 77-1013

AVIATION RESEARCH LABORATORY

INSTITUTE OF AVIATION
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

(12)

TECHNICAL REPORT

ARL

AN EVALUATION OF A COMPLEX COMPUTER-BASED FLIGHT PROCEDURES TRAINER

STANLEY R. TROLLIP

AD No.
DDC FILE COPY

ARL-77-1/AFOSR-77-1
JANUARY 1977

Approved for public release;
distribution unlimited.

PREPARED FOR

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFSC)
NOTICE OF TRANSMITTAL TO DDC
This technical report has been reviewed and is
approved for public release IAW AFR 190-12 (7b).
Distribution is unlimited.
A. D. BLOSE
Technical Information Officer

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER 19 18 AFOSR-TR- 77-1013	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER 9
4. TITLE (and Subtitle) 6 AN EVALUATION OF A COMPLEX COMPUTER-BASED FLIGHT PROCEDURES TRAINER		5. TYPE OF REPORT & PERIOD COVERED Interim Scientific Report
		6. PERFORMING ORG. REPORT NUMBER 14 ARL-77-1/AFOSR-77-1
7. AUTHOR(s) 10 Stanley Raynes Trollip		8. CONTRACT OR GRANT NUMBER(s) 15 F44620-76-C-0009
9. PERFORMING ORGANIZATION NAME AND ADDRESS Aviation Research Laboratory Institute of Aviation University of Illinois, Savoy, IL 61874		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 16 61102F 2313/A4 17 A4
11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Office of Scientific Research (NL) Building 410 Bolling AFB, DC 20332		12. REPORT DATE 11 January 1977
		13. NUMBER OF PAGES 146
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 12 152p.		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release: distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Computer-based trainer	Simulation
Flight procedures trainer	Holding pattern
Flight training	Human Factors

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A computer program was written for the PLATO IV computer-assisted instruction system to teach pilots how to fly holding patterns in various wind conditions. The students participated actively in the learning process by "flying" a simulated airplane by means of a hand-controller which was interfaced with the computer system. Fifty-two subjects were used to evaluate the program in terms of how well it taught, how well the strategies taught would be applied to different situations, and how the subjects reacted to this mode of learning.

1B

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

406 171

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. The subjects were divided into four groups, two control and two experimental. The control groups were given a lecture on how to fly holding patterns, while the experimental groups completed the computer program. One control and one experimental group were trained to criterion in a ground trainer how to fly holding patterns in different wind conditions. On completing this the two groups had to fly one procedure turn in a difficult wind condition. This maneuver requires the application of strategies similar to those used in flying holding patterns. The second control and second experimental groups flew the procedure turn first, followed by the holding pattern sequence.

The results on the holding pattern sequence showed the experimental groups to perform better on all measures, with significant differences ($p < .02$) between groups on the number of critical errors to criterion. However, there were no significant differences between groups on the performance of the procedure turn. Most subjects responded very positively to the PLATO instruction and many commented that being able to "fly" the patterns and then see a display of their flight paths facilitated learning. The results also indicated that the experimental subjects prepared themselves better before each pattern, and were better oriented when flying it.

The results of the evaluation suggest that many procedures may be taught better using an inexpensive large-scale computer system which simulates the necessary cues rather than a trainer which simulates the physical environment.

AVIATION RESEARCH LABORATORY

Institute of Aviation

University of Illinois at Urbana-Champaign
Willard Airport
Savoy, Illinois
61874

Technical Report

ARL-77-1/AFOSR-77-1

January 1977

AN EVALUATION OF A COMPLEX COMPUTER-BASED FLIGHT PROCEDURES TRAINER

Stanley Raynes Trollip

Prepared for

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE

406 171

B

TABLE OF CONTENTS

1. INTRODUCTION

The PLATO system (Alpert and Bitzer, 1970) at the University of Illinois is a general purpose computer-assisted instructional (CAI) system. It differs from other existing systems, however, in a number of ways. Unlike those that have been used for several years, such as the IBM 1500 series which were adaptations of existing equipment, PLATO was conceived, designed and built as a teaching device. Experts from the education, economic and engineering fields collaborated in its design and have monitored its development through four generations. It is both large-scale and user-oriented, and designed to operate 1000 terminals simultaneously with a maximum response time of no more than half a second.

The recent introduction of large scale CAI systems like PLATO has resulted in a proliferation of different types of lesson materials that could not have been anticipated a decade ago when the prototypes of present systems were being implemented. It is not so much the diversity of subject matter (see Lyman, 1974) that was unanticipated, but rather the variety of instructional techniques and sophisticated applications. Unlike more primitive systems which have been restricted in general to drill and practice or testing situations, PLATO, with its powerful graphic and computational capabilities, permits a bewildering spectrum of applications.

One programming group has an extensive library of programs designed to teach young children how to read. To accomplish this, random-access

audio devices are used to 'speak' to the child, and touch-sensitive panels sense where the child touches the terminal's display screen. Another group has designed lessons in which the student must diagnose the illnesses of animals by asking PLATO such questions as: "What is its respiration like?" or "What does its tongue look like?" or "I would like to do a urinalysis. What results would it give?" A third group has simulated a chemistry laboratory, another teaches Chinese or Russian, another writes programs to teach technical courses at the Air Force Academy. Still another group is investigating the use of PLATO as a computer-based communications medium to implement such activities as "computer conferencing," which, in the long run, could produce significant educational effects by allowing students to "audit" current scientific debates between practicing scientists.

There are lessons available to teach users how to use PLATO and its programming language, TUTOR. There are service programs to aid in the design of complicated graphics, (hence, even the programming is automated to a small extent) and there are programs providing a variety of statistical analyses and tests. Programs characterized by different instructional techniques and a variety of additional external output devices, such as the audio-device, the slide-selector, or music box, constitute the bulk of programs. However, there are a few programs, such as those written by the reading group, that use external input devices.

The most widely used external input device, other than the keyboard, is the touch panel, which is fitted over the screen

of the terminal. Touching the screen with a finger causes an interrupt in a 16 x 16 grid of infra-red beams and allows the computer to "sense" where the screen was touched. In other applications external input facilities are used to monitor the design of electrical circuits (Neal, 1974), feedback then being provided appropriate to the problem.

Analog input capability has also been made possible by using an analog/digital interface (Battle, 1974), which was designed specially for the application discussed in this dissertation. Both automobile (Parker and Voss, 1975) and airplane controls (Trollip, 1975) have been linked through this interface to the system for use in instructional programs.

In the light of these developments one can say that CAI is coming of age. From its tentative beginnings less than 20 years ago, it has grown in applicability and availability to such an extent that credit courses from primary school through college level are taught regularly at many locations around the country.

However, despite widespread research and development, a certain conservatism dominates the field -- a conservatism resulting largely from a lack of sufficiently flexible systems. Unfortunately this conservatism tends to prevail even on systems like PLATO that no longer dictate the type of instruction available. Much of what is taught is taught as a series of rules or principles, and most CAI has capitalized on the fact that computers are well-suited to the branched presentation of material of this kind. Consequently little research exists into the uses of CAI in areas such as skill acquisition where the emphasis would be "knowledge how" rather than "knowledge that."

The research discussed in this dissertation evaluates a PLATO program (Trollip, 1975) whose purpose is to teach "how" rather than "that." The program simulates interaction with the real world by allowing pilots "to fly" an "airplane" through certain maneuvers by means of a hand-controller or "joystick" interfaced with the PLATO system.

2. BACKGROUND

The rationale for the development of a PLATO program to teach pilots instrument-referenced maneuvers was rooted in three different areas. First, the project was undertaken because it was felt that it would be cost effective, competing successfully with alternate forms of instruction. Second, the "learning by doing" philosophy as expounded by people such as Bruner (1973) and Papert (see Papert and Solomon, 1972; Papert, 1972a; 1972b) could be incorporated into such a training program. Third, initial efforts in the uses of imagery techniques in training as reported by Prather (1973) could be extended and their effectiveness investigated.

At present, the training of complex skills particularly in aviation, is an expensive process. In the field of general aviation, where the aircraft used are usually less expensive than their Air Force counterparts, the cost of an advanced rating can be very high. At present, at the University of Illinois, which boasts one of the least expensive curricula, the cost of the instrument rating courses alone is $1400. From a fixed-base operator, this cost is generally about 50% higher due to the fact that the student has to pay for the instructor as well as the plane, and because a profit margin is also built into the cost. Because of the high cost of training, more and more extensive use of flight simulators is being made. However, even these are expensive, with the purchase price being about the same as the plane being simulated. The benefits, of course, come in the elimination of accident risk and the

reduced operating cost per hour. A single-engine plane may rent from $15 to $30 per hour, while an appropriate simulator will rent from $5 to $15 per hour depending on whether the simulator also provides motion simulation. Whether or not motion simulation improves transfer of learning to the airplane is a question of considerable controversy with ardent exponents on both sides (see Jacobs and Roscoe, 1975). This is discussed further in the conclusion.

In either the plane or the simulator, some considerable portion of the time is used, not by practicing those procedures that need practice, but rather by completing the mundane and relatively unimportant procedures like starting the engine, climbing to altitude, etc., which do not require extensive rehearsal. For greater efficiency both in terms of time and money, CAI programs like the one discussed here have a great deal of potential. Not only is the operation of such programs substantially less costly than even a simulator (less than $1.00 per student contact hour for PLATO), but there is also a reduced need for an instructor because the program can offer feedback frequently comparable to that of an instructor. In a group setting, one instructor can monitor the ongoing performance of several students at once on a master terminal.

However, the fact that the student no longer has to do all the things involved in flying the airplane, such as communications, navigation, and clearance, has greater implications than merely saving time and money, for it allows him to provide full attention to the task at hand, rather than have his attention divided by these tasks. After all, the purpose of the program is not to teach a person how to fly a plane,

but to teach him to understand how to fly instrument maneuvers. When these procedures have been mastered, then it is time to put them into practice. This isolation of the pertinent task is one of the powerful capabilities of CAI systems in general and has been used with considerable success in the teaching of subjects such as chemistry, physics and nursing to name but a few.

For example, in one Chemistry curriculum students use CAI lessons as an introduction to each session in the laboratory (Smith, 1970; Smith, Ghesquiere, Avner, 1974). In this way, by simulating the necessary laboratory equipment such as pipettes, burettes, and chemicals, the student has the opportunity to complete the entire laboratory assignment without ever having to handle the real equipment. Once he has successfully completed the PLATO assignment, he can then go into the laboratory and perform the experiments knowing what is to be done. Time in the laboratory is diminished, much confusion is eliminated, and the student learns to prepare himself before conducting an experiment. Furthermore, because the PLATO laboratory is safe to use, it is possible for a student to make exploratory tests with complete safety.

In the same way, the instrument program allows the student to understand the dynamics of instrument flight and to practice what he has learned in a non-threatening environment with the computer giving appropriate feedback. Thus the student knows what it is all about when he gets into the plane and is asked to perform the maneuver. He does not have to concentrate on keeping the plane flying and at the same time to work out appropriate procedures. Most attention can be given to the task of keeping the plane flying safely.

This use of CAI represents a dramatic change from the traditional application. Normally CAI lessons have taught by merely providing the necessary information, sometimes in a manner that could not be done either by a teacher or a book. Anything learned would be as a result of the student's digesting the information and remembering it. On the other hand, the instrument program teaches by allowing the student to actually experience and do what is to be learned. On completing the training sequence, the student has learned far more than the laws and rules governing the maneuvers. He has had the opportunity to develop and use these rules, and test them in an operational situation. In this way, the student learns the dynamics of what is to be done by doing it rather than by trying to project to the real situation. Looked at like this, the PLATO instrument program actually simulates a simulator, whereas a traditional program would simulate a book or perhaps a teacher.

One other technique incorporated into the program was that of using imagery to improve performance. Prather (1973), for example, used a type of imagery instruction or mental practice as an adjunct to normal training to teach Air Force Academy students how to land a T-37 airplane. In the study Prather played voice recordings of where the plane was in the landing pattern and what should be done to the controls moment-to-moment. Over the training sessions, the amount of detail in the tapes diminished. So in the first session the instructions were comprehensive, including airspeeds, throttle settings, altitudes, etc. However, in the final session the

recording would merely be, "You are now on downwind" and "You are now on final." The students sat in a cockpit mock-up and had a throttle and a control stick to manipulate but no working instruments to read. What was being taught was a mental representation of the landing procedure, so that the students knew what to do by referencing this representation. This is similar to the subjective experience of "keeping ahead of the plane." In the evaluation of Prather's techniques, the experimental group landed better than their control counterparts, who learned conventionally.

A similar technique was used by Feurzeig (1971) in two computer-based applications. The first, which provided much of the inspiration for this dissertation, was a program to teach pilots to fly holding patterns. The second was designed to teach the skills involved in estimating the relative courses of ships moving at different rates and in various configurations. In the evaluation of their programs Feurzeig and his associates found that the instructional programs were effective.

A more abstract approach based both on the notion that the student should have an appropriate mental picture of the task, and that this was best obtained by making the student an active participant in the learning process was suggested by Goldstein and Goldstein (1972) They made use of the TURTLE in the LOGO project (see Ableson, Goodman, and Rudolph, 1974) to let prospective pilots program different situations related to flying. "The goal is to provide a better environment for a pilot to build mental models of his plane's performance under different

flight conditions and graphically explore situations he could never safely be exposed to in the air." (Goldstein and Goldstein, 1972, p. 3). Making the student actually program the TURTLE to "fly" in different wind situations increases the probability that he will acquire a complete understanding of the concepts involved, while the output of the programs allows him to examine a picture of what happens under the different conditions. In this way the necessary abilities for adequate "thinking ahead" of the plane are provided.

The examples described above indicate that the use of this type of imagery training could lead to improved performance in two ways. First, the pilot would be provided with a readily accessible and vivid reference as to his performance if he kept in his "mind's eye" a picture of where his plane was in the pattern and how it was progressing relative to the ground. Thus, if his attention were temporarily distracted by having to compensate for some error, it would be easier for him to regain orientation. The second improvement of performance should be apparent in the pilot's preplanning abilities. If the task to be performed were mentally visualized, then, perhaps this would of itself cause the pilot to be "thinking ahead" of the plane, an important ingredient for safe flying. By being prepared for various contingencies, the shock of surprise at an unexpected occurrence would be minimized, and the appropriate corrective actions implemented quicker.

It was for all these reasons that a program was written (Trollip, 1975; Trollip and Ortony, in press) to teach pilots how to fly holding patterns -- a traditionally difficult and expensive problem: difficult, because wind necessitates changes in the shape of the pattern to

compensate for drift; expensive, because suitable experience has hitherto only been possible through the use of either a simulator or a plane. The program makes use of a hand-controller or "joystick" interfaced into the PLATO system, the purpose of which is to provide the student with controls similar to those in an airplane.

3. THE PROBLEM: TEACHING PILOTS TO FLY HOLDING PATTERNS

The PLATO program used in this research simulates a complex training task, the purpose of which is to teach pilots effective strategies for flying holding patterns in varying wind conditions. A holding pattern is a maneuver designed to ensure that the ground controllers know exactly where each plane is in a high density traffic situation to facilitate sequencing and spacing of landings. Orders to hold are usually given when a number of planes are waiting to land at the same airport or when a plane has to await clearance to proceed on its course. It is a means of stacking many planes in a small area while ensuring maximum safety.

It is generally recognized that holding pattern flying is one of the most difficult procedures to teach even though it is conceptually easy and can be worked through with a pencil and paper. The major problems occur when there is wind which necessitates a change in the shape of the pattern to compensate for resulting drift. Most instructors would agree that insufficient preplanning is the cause of most of these problems, and it is the teaching of this planning that is the overt purpose of the PLATO program.

In a no-wind condition the holding pattern is racetrack in shape (see Figure 1) with the end of one straight leg defined by some radio fix. This fix may be a radio transmitter or the intersection of two defined radials broadcast by transmitters some distance away. The pattern is flown with the radio fix marking the end of the inbound leg. There follows a 180° standard rate turn (3° per second), to the right unless

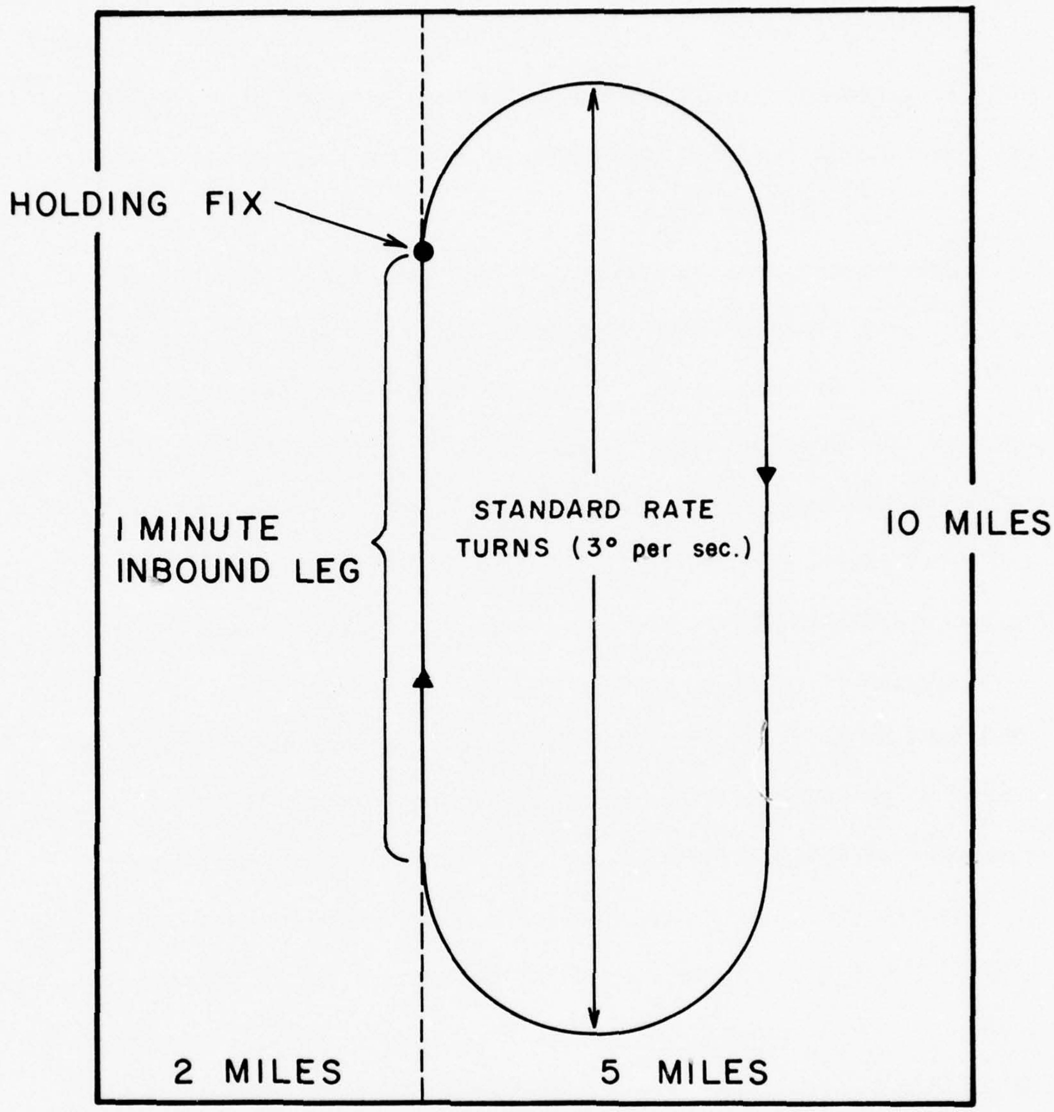

Figure 1. Shape of holding pattern in a no-wind situation, as seen from above.

otherwise specified. The size of the pattern is defined by making the inbound leg one minute long, and the direction is defined by specifying the course of the inbound leg. To ensure adequate safety each aircraft is assigned protected airspace which is five miles wide on the holding side of the inbound leg, two miles wide on the non-holding side, and ten miles long.

In addition to these specifications, the altitude at which the pattern is to be flown is specified, with 1000 feet vertical separation between aircraft allowing many planes to hold over the same area on the ground. If, for example, these planes were all waiting to land at the same airport, the plane lowest in the stack would be cleared to land and all the others above it cleared to the holding altitude below. New arrivals would be stacked at the top. Figure 2 shows a side-on view of the holding stack.

The task of flying holding patterns would be relatively simple if there were never any wind. Unfortunately this is seldom the case, and consequently, pilots need to be taught procedures to deal with windy conditions. As has already been mentioned, one of the determining features of the holding pattern is that the inbound leg be one minute long, and it is this feature that remains constant over all situations. If the inbound leg, whose direction is specified, is to remain one minute long irrespective of the wind, then the shape of the pattern must change to accomplish this.

For example, if a plane were flying at an airspeed of 80 miles per hour, with a 20 mph headwind on the inbound leg, the ground distance

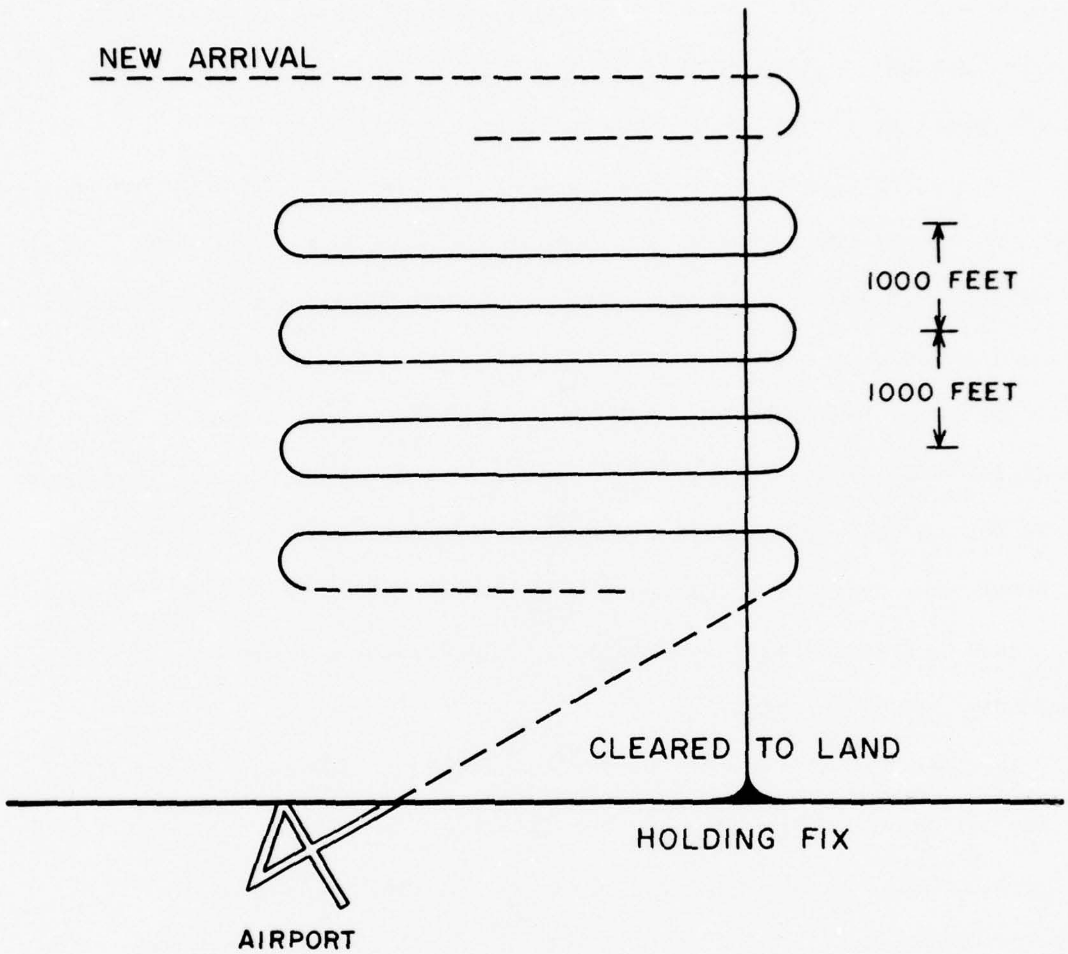

Figure 2. Side view of the holding stack.

covered in the one minute on the inbound leg would be less than if there were no wind. In fact the length of the inbound leg would be exactly 1 mile (because the groundspeed is now 60 mph), as opposed to 4/3 mile in the windless situation. This can be seen from Figure 3.

Continuing the example, because there was a 20 mph headwind on the inbound leg, on the outbound leg there would be a 20 mph tailwind resulting in a groundspeed of 100 mph. Because the length of the outbound leg would be approximately equal to that of the inbound, it would be about 1 mile long. At 100 mph the time on the outbound leg would be (1x60)/100 minutes, or 36 seconds. Thus one procedure that is available to deal with different headwind components is to estimate the outbound time to give an inbound leg of one minute. Initially this time is estimated from knowledge of existing wind conditions and is improved each time around the pattern.

In crosswind situations the difference in groundspeeds occurs not on the inbound or outbound legs, but rather in the turns. When turning into the wind, the groundspeed drops, and because the rate of turn remains constant at 3^{o} per second, the radius of the turn decreases. When turning downwind, the groundspeed increases resulting in a turn of larger radius. In addition, in a crosswind situation it is necessary to head the plane into the wind on both the inbound and outbound legs to maintain the desired course. This is called 'crabbing' into the wind. Figure 4 illustrates what the pattern shape looks like if there is a strong wind from the left on the inbound leg.

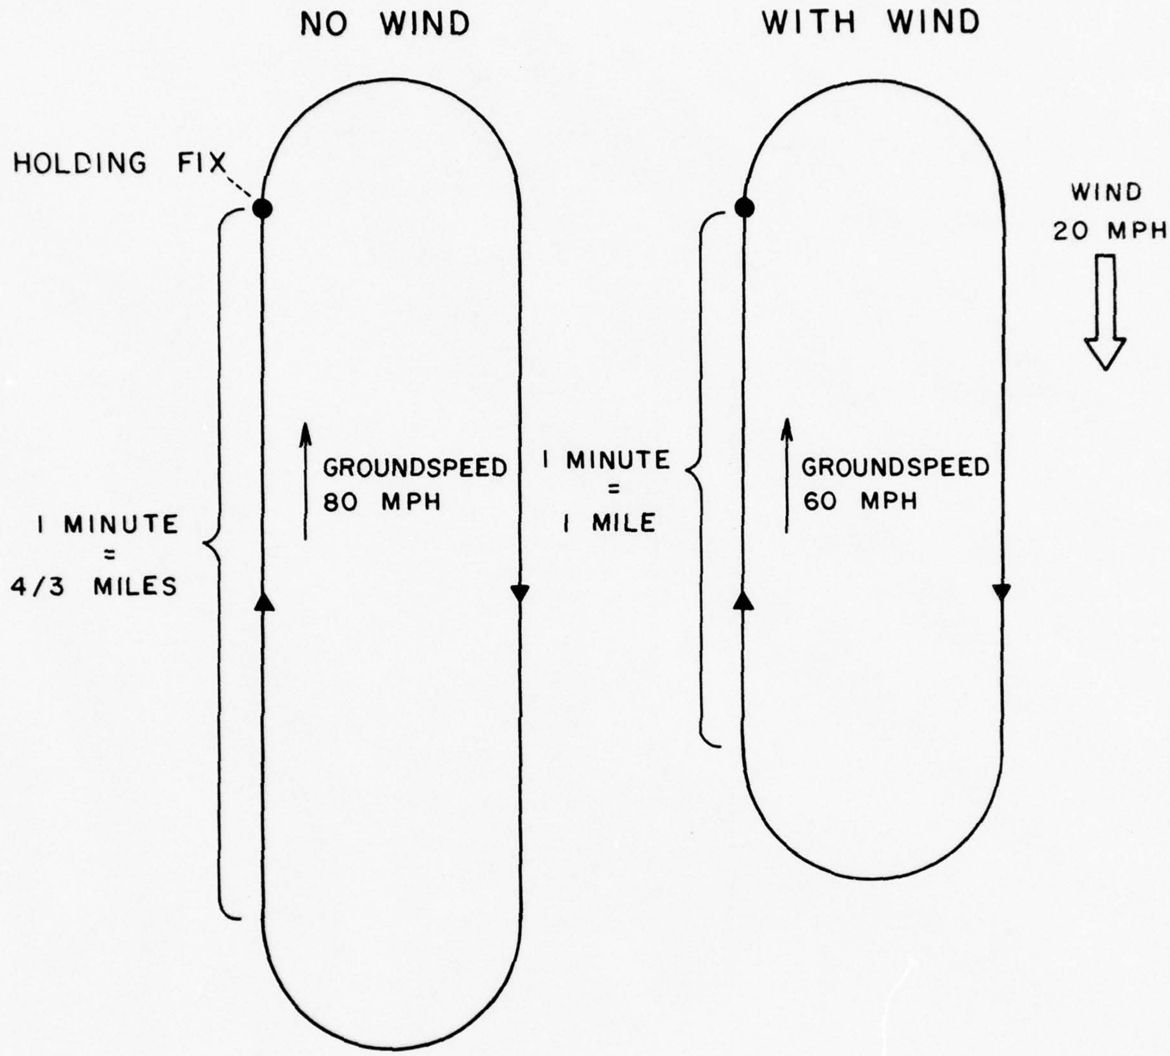

Figure 3. Shape of pattern in no-wind and headwind situation.

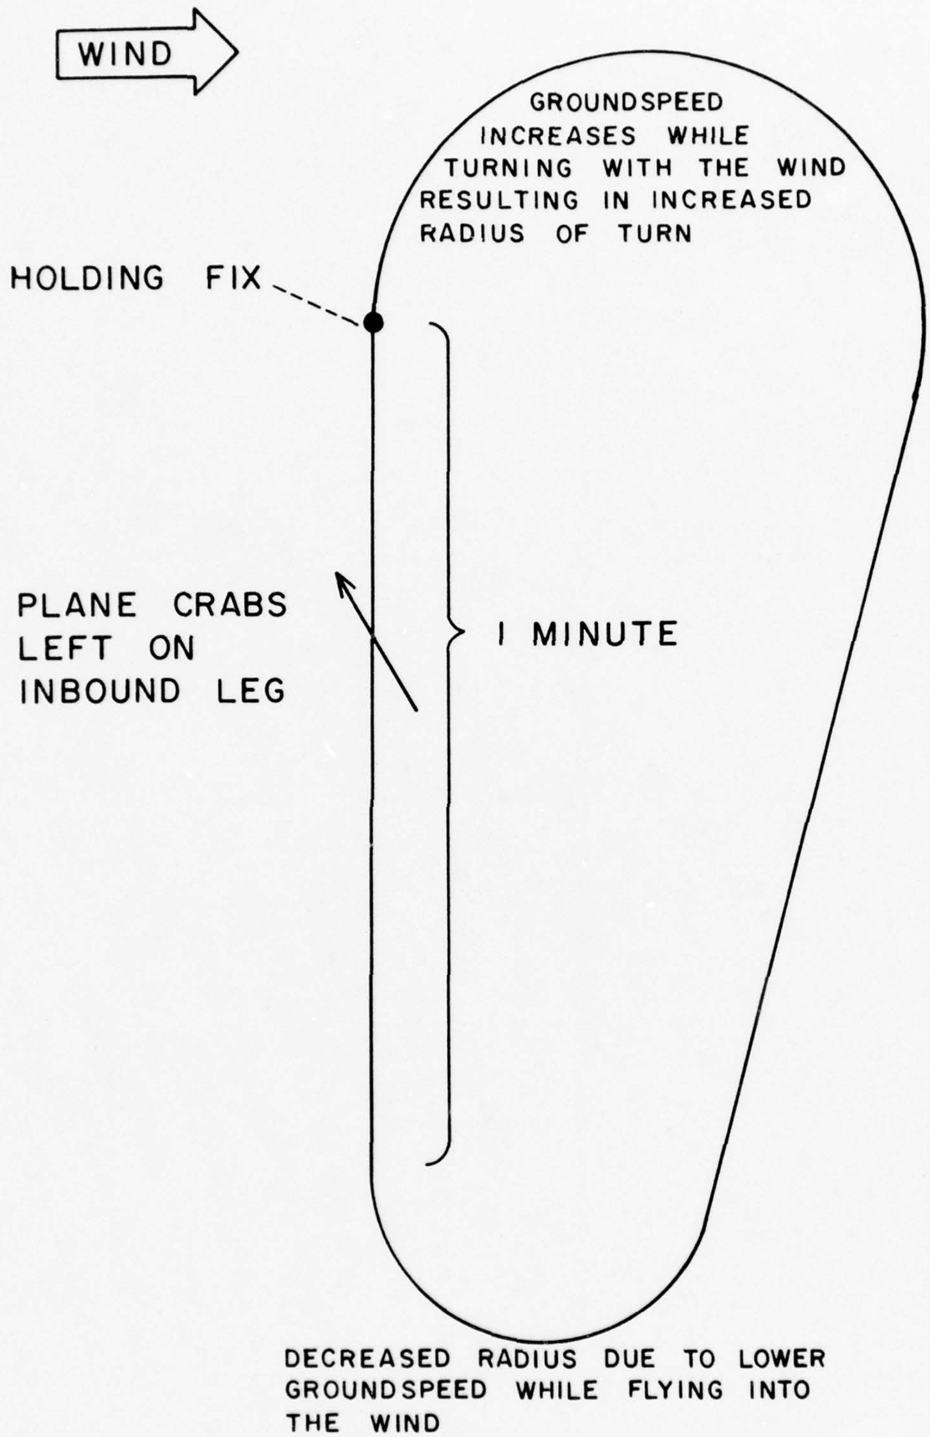

Figure 4. Shape of the holding pattern in a typical crosswind situation.

The important question for the pilot is how to determine the outbound heading in order to join the top and bottom turns which are of different sizes. The rule of thumb method is to hold twice the inbound crab angle on the outbound leg. Thus if a 10^{o} crab held course on the inbound leg, then holding a 20^{o} crab on the outbound should join the two turns. The rationale behind this heuristic is that one cannot crab while turning; therefore by holding twice the crab angle on the outbound leg one compensates for the lack of crabbing in the turns. In practice this works quite well.

In the usual wind situation where the wind is neither a headwind nor a direct crosswind, integrating both correction methods is the desirable solution to the problem.

As can be seen from the foregoing discussion, the techniques for compensating for different wind conditions are not excessively complicated. However, it is the experience of most flight instructors that performance on this maneuver is not good. Perhaps one reason is that traditionally pilots are taught that the holding pattern is racetrack in shape with minor variations for different wind situations. What should probably be taught is that in most situations the holding pattern is <u>not</u> shaped like a racetrack, but is distorted according to the wind direction and strength. That is, emphasis should be placed on how and why the shape of the pattern should change to ensure that the inbound leg remains one minute long. By learning the task in this way the pilot should not only have a better understanding of the relationship of the plane to the ground in terms of speed and direction, but he should also be able to visualize

better where in the pattern he should be and in which direction he should be flying. The PLATO program evaluated in this dissertation approaches the problem in this way, emphasizing the dynamics of the pattern, and its relationship to the wind.

4. THE SOLUTION: THE PLATO PROGRAM

To accomplish its purpose the PLATO program is divided into two distinct parts. The first provides the necessary information regarding the details of the holding pattern, such as its definition and the clearance given by the ground controllers. The second allows the pilot to "fly" holding patterns using the hand-controller interfaced with the system, and the instruments simulated on PLATO's screen.

In the belief that CAI should not be used redundantly, nor should it provide instruction available more readily or more inexpensively elsewhere, the didactic part of the PLATO instruction is structured as follows: on arrival at the PLATO terminal for the first time, the student is given a typed booklet outlining all the information necessary to fly holding patterns. This information includes definitions, and strategies how to compensate for wind-induced problems. By providing the bulk of the information in this way, one does not tie up the system with a "page-turner" program.

When he has finished reading the booklet the student signs onto the PLATO terminal and is taken directly to the first part of the holding pattern program. In this part the program quizzes the student about what he has just read, requiring the student to make the same decisions that he would when cleared to hold in the real situation.

First, the student must decide where the pattern for which he has been cleared, will be flown in respect to direction of the inbound leg and

direction of the turns. For each clearance given, there appears on the PLATO screen a diagram of the holding fix and environs. By touching the screen in the appropriate position the student can indicate exactly where the pattern will lie in relation to the fix. Having a clear orientational picture of the hold prior to arrival at the fix facilitates performance and gives the pilot time to make other decisions.

After successfully completing three of these questions, the student is given a series of clearances with typical wind information such as the Winds Aloft Forecast or the reported surface winds in the area. In reality neither of these wind reports is accurate for the holding altitude, but rather each gives the pilot an indication of what winds may be expected at the holding fix. From this information the student is required to make an estimate of whether the outbound leg should be timed at less than or more than one minute for the inbound leg to be the prerequisite one minute. In the real situation it is good practice to have made this decision before arriving at the fix for the first time.

Finally, the program allows the student to experiment with different wind directions and speeds to see how the ideal pattern will change its shape. The student types in a wind direction and speed, and PLATO responds by drawing a picture of what the ideal holding pattern should look like with respect to the ground. PLATO then superimposes the track of an airplane flying in the same wind conditions without compensating for the wind. The magnitude of the deviation from the "traditional" race-track shape surprised even some experienced pilots, and the difference between

the ideal pattern and the one flown without compensation made the point very well that some strategy was essential even on the first time around. Figure 5 shows the ideal pattern for a strong headwind (top) and for a strong quartering tailwind from the right (bottom). Also shown in each case is the path of the plane flown as though there were no wind at all.

The second part of the PLATO program allows the student to "fly" holding patterns in different wind conditions using the interfaced hand-controller (Figure 6), and the appropriate instruments simulated on the screen (Figure 7). Although Figure 7 shows an altimeter and a rate of climb indicator (on the left), experience with pilots before the study indicated that the inclusion of altitude control added nothing to the program (except perhaps unnecessary distraction), and consequently throughout the program the altitude was kept constant.

After being given adequate opportunity to become familiar with the operation of the hand-controller, the student is required to fly a series of holding patterns starting with the simple no wind situation, progressing in difficulty to a direct crosswind and finally to a strong quartering headwind condition. In each case the student has to satisfy the program's "instructor" that performance was adequate before being allowed to proceed to the next pattern.

At the end of each pattern the program draws on the screen a picture of how the pattern was flown and superimposes the ideal pattern for that wind condition. From information collected while the pattern was being flown, an assessment is made (see Chapter 6 for details) as to how the student performed. Feedback is also provided on all stages of the pattern.

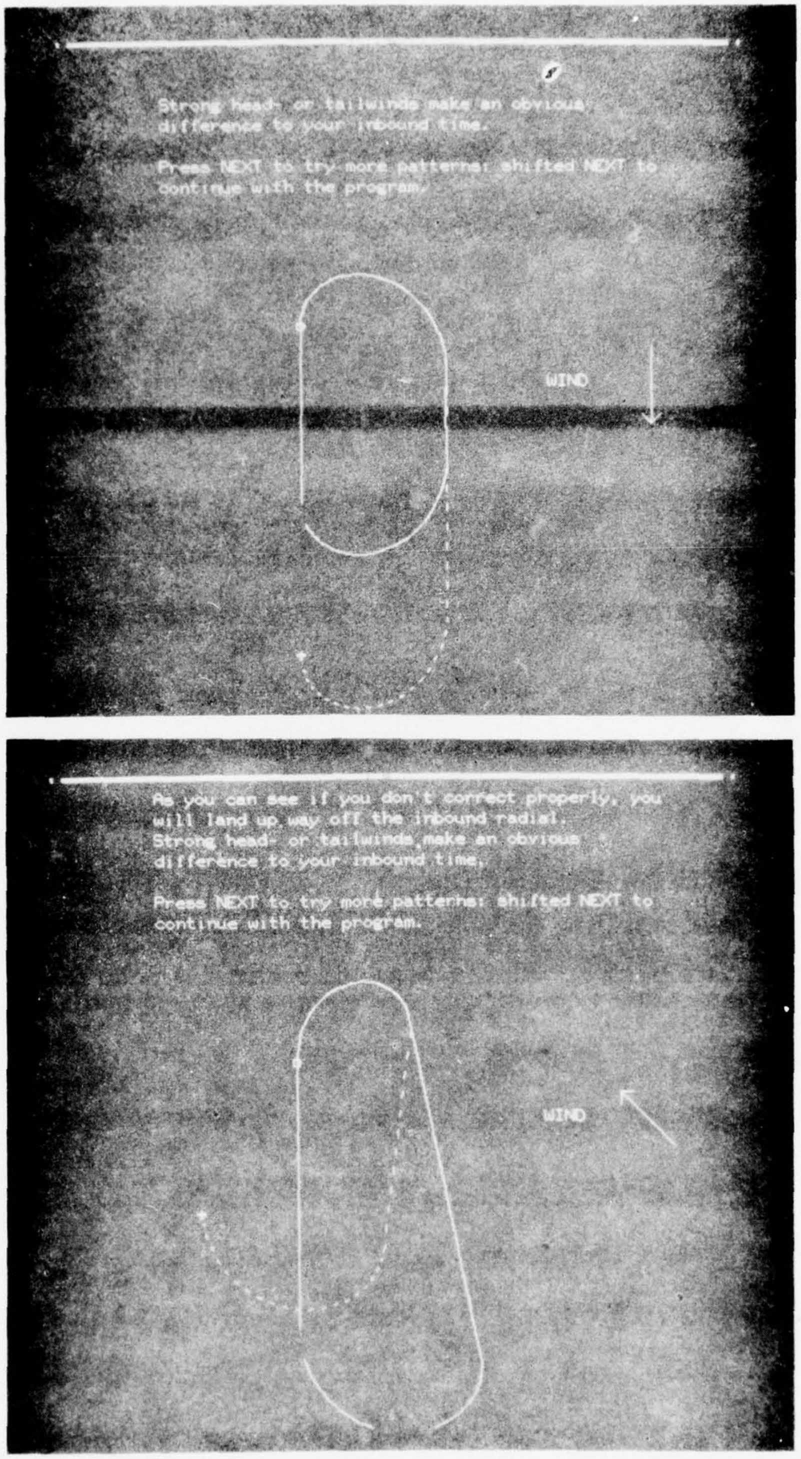

Figure 5. PLATO output showing the ideal pattern (solid line) and uncompensated track (dotted line) for different wind conditions.

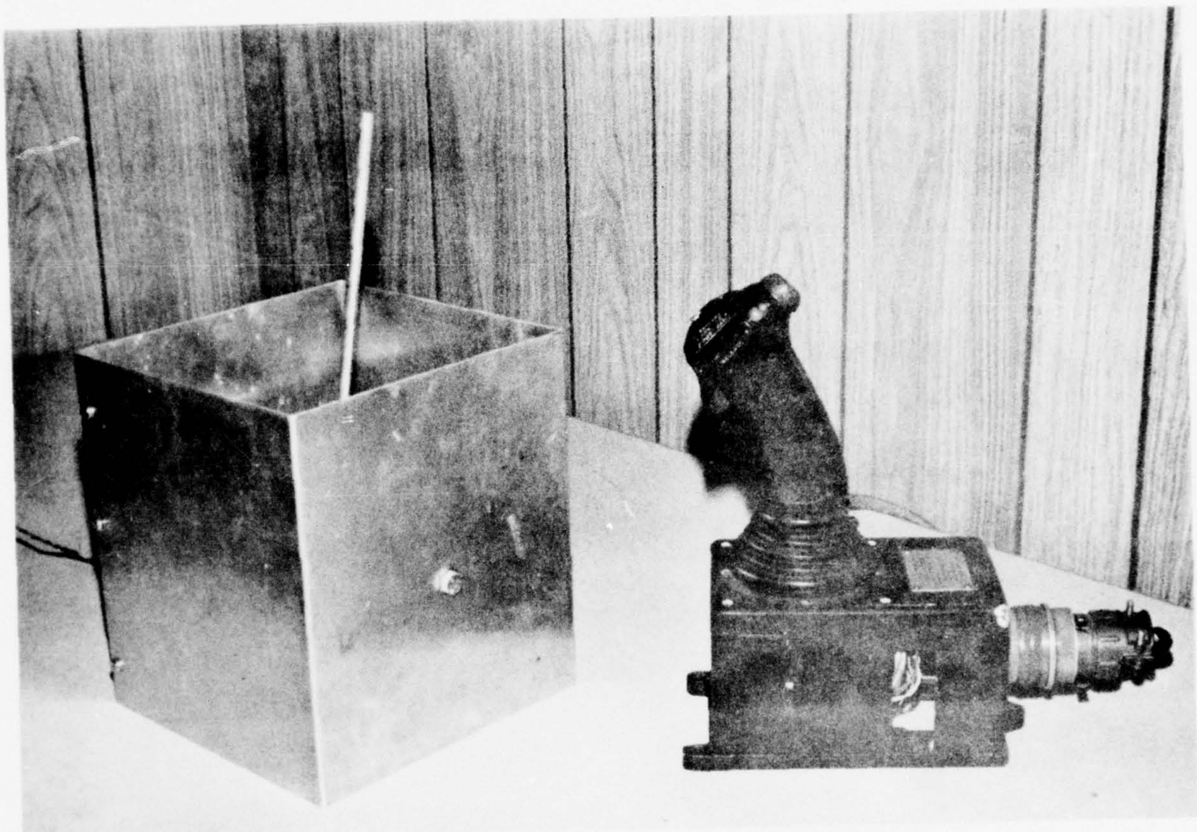

Figure 6. The Inexpensive home-built controller (left) worked just as well as its expensive commercial counterpart.

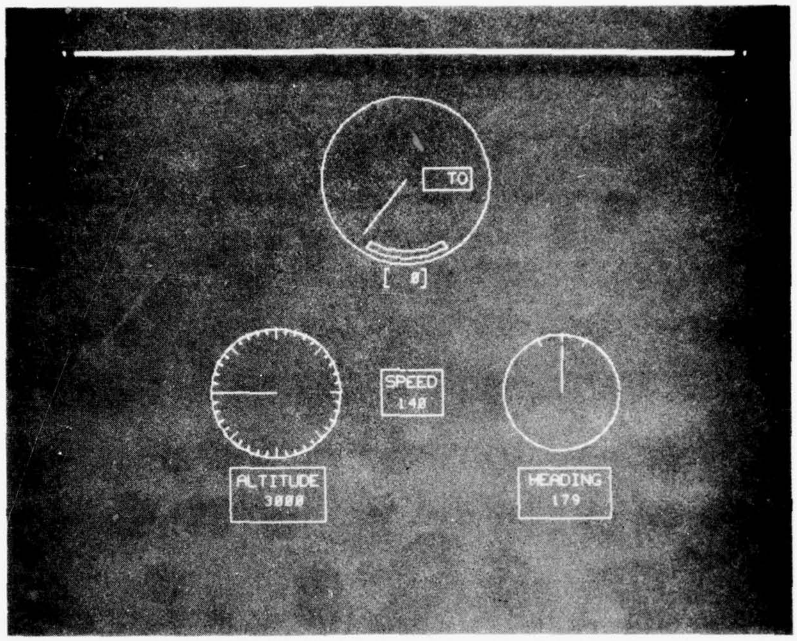

Figure 7. The PLATO instrument panel.

Figure 8 shows a pattern flown in the quartering headwind situation where the crab on the outbound was not sufficient. Below the picture of the PLATO screen are all the comments made about the resulting pattern by the program. As can be seen, although much of the feedback is of the nature of pointing out mistakes, the program also attempts to suggest causes and remedies for these errors.

Thus on completing the PLATO program each student has had the opportunity to learn the facts about holding patterns and to make the necessary decisions regarding their spatial orientation and the strategies to employ for their successful execution. Furthermore, each student has had the opportunity to "fly" various holding patterns in a representative sample of wind conditions. He can see immediately how well the strategy he used worked, and is given appropriate feedback from an "instructor."

Because the student does not have to fly a plane to accomplish all this, all his attention can be focused on the cognitive aspects of completing the pattern, especially the formulation of the correct strategies. Thus, at least in theory, when the student has to fly holding patterns in a plane, his attention can for the most part be directed at flying safely rather than at searching for the correct procedures to accomplish the pattern.

The PLATO program, therefore, does not approach the task of teaching holding patterns traditionally, but emphasizes the changing nature of the pattern and the effects of wind on pattern shape. It also uses an automated interactive CAI simulation which requires the student to participate actively in the learning process.

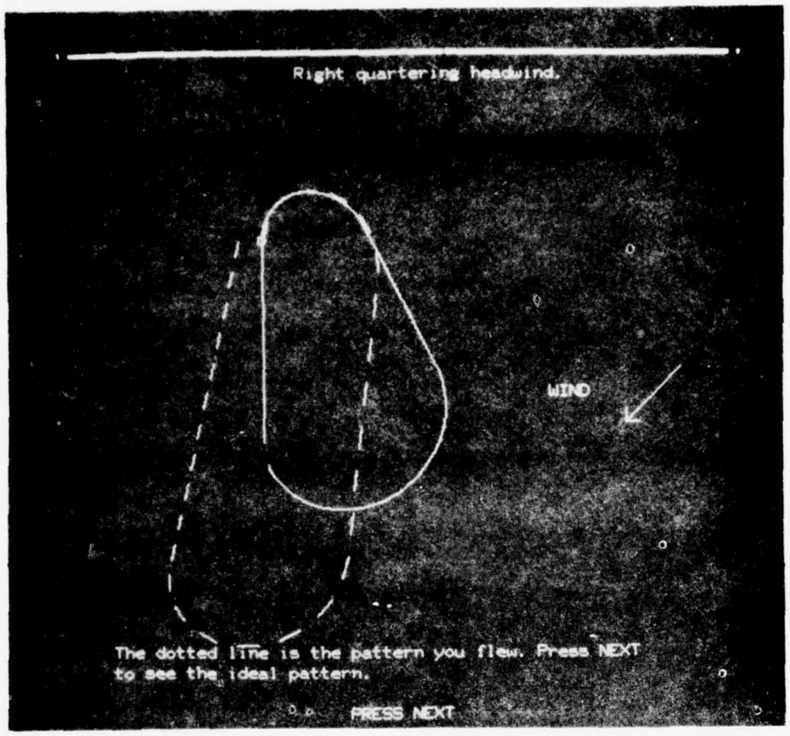

Figure 8. Student compensated incorrectly for the wind. Photograph shows graphic feedback.

5. THE ISSUES

The complete evaluation of the holding pattern project covers many different aspects which fall into two main categories: the theoretical issues and the evaluative issues. These may be summarized as follows:

<u>Theoretical Issues</u>

1. A scrutiny of the interaction of the holding pattern program and the PLATO system in terms of both the hardware and software;
2. The establishment of procedures for evaluating performance on a continuous task.

<u>Evaluative Issues</u>

1. The determination of the instructional effectiveness of the holding pattern program in training pilots to fly holding patterns;
2. The cost effectiveness of such training;
3. The determination of the generalizability of the instruction given in the PLATO program on holding patterns as measured by performance on flying procedure turns, a related but distinct task;
4. The reaction of the subjects to the PLATO program.

For the most part the theoretical issues mentioned above are discussed in Trollip (1975) where it is concluded that the holding pattern program and the PLATO system do not interact with any significant detrimental effects. Since that report, further improvements have been made to the hardware,

expecially the analog/digital interface, allowing simpler and more efficient software programming. These modifications are elaborated in Battle (1974).

Trollip (1975) also discusses a proposed method to be implemented in the PLATO program to evaluate how well the student flies each pattern. Attempted empirical validation of this method found problems sufficiently serious to cause it to be abandoned. The details of this validation and a description of the method of evaluation finally used are discussed in the following chapter.

Instructional and Cost Effectiveness

The major issue at hand is, of course, whether the PLATO program meets its instructional objectives and if so, does so in an economical way.

The approach taken in this dissertation is to compare these instructional objectives with the situation that one would normally find at most fixed-base operators. That situation is where an instrument student wanting to learn how to fly holding patterns would be given one or more groundschool lectures on the theory of the patterns, followed by instruction in an airplane or simulator. The hypothesis to be tested was that the PLATO-based training sequence could replace the ground school instructor completely and, at the same time, provide a learning experience sufficiently better that the flight time required to achieve a predetermined standard would be reduced and, finally, that fewer serious errors would be made in attaining that criterion.

Although this approach makes the determination of the program's cost effectiveness very difficult, due to an inability to put a price on the original groundschool instruction, it is the most valuable initial approach because it allows an evaluation of the program on a stand-alone bases. It is not the purpose of this first evaluation to determine the best pedagogical approach to the teaching of holding patterns, but rather to determine whether the PLATO program can do the job at all, and whether further evaluations are called for.

Consequently, use of some of the traditional evaluative tools, such as the transfer effectiveness ratio, TER, (Roscoe, 1971; Povenmire and Roscoe, 1973) cannot be made, due to the formative nature of the evaluation.

Generalizability of Instruction

Because part of the appeal of a computer-based simulator as described throughout this dissertation is its flexibility, it is necessary to evaluate not only how well pilots fly holding patterns, having been taught by the PLATO program, but also how well they can generalize that learning to similar but different situations. The "other" task chosen for this study was that of flying procedure turns, a maneuver that requires the application of strategies akin to those of holding patterns to compensate for the effects of wind.

A procedure turn is a maneuver used in making instrument approaches to landing, whereby the airplane reverses its direction of flight while minimizing the chance of deviating laterally too far from its prescribed course. Typically, the plane flies away from the airport on a specified radial of a navigation transmitter until it passes the outer marker (OM),

an electronically recognizable point some miles from the airport. Shortly after crossing the OM outbound, the plane has to make a left (sometimes right) turn of no more than 90^{o}, followed shortly, thereafter, by a 180^{o} standard rate turn. The plane should now be flying towards the prescribed radial, intercepting it at a reasonable angle. As the navigation instruments begin to indicate that the radial is being approached, the pilot would initiate a turn so as to end up flying toward the airport on the same radial as it flew from it. Figure 9 depicts a typical procedure turn.

As in holding patterns, the strength and direction of the wind determine for how long and on what headings the pilot should fly the legs of the turn. Insufficient compensation for wind regularly causes the plane to be blown across the inbound radial, or even inside the OM before being established on course. In either case, efforts to reestablish the correct position could well lead to potentially disasterous results such as the pilot forgetting to extend the landing gear, or maintain a safe altitude.

Acceptability

An important part of any evaluation of a new instructional technique (or medium) is the assessment of its acceptability to the user. Anecdotal evidence exists of a statistics class which was divided into control and experimental groups. The experimental group was taught statistics by PLATO III (the predecessor to the current PLATO IV), while the control group, but none of the experimental subjects signed up for a second statistics course, whereas several of the control group did.

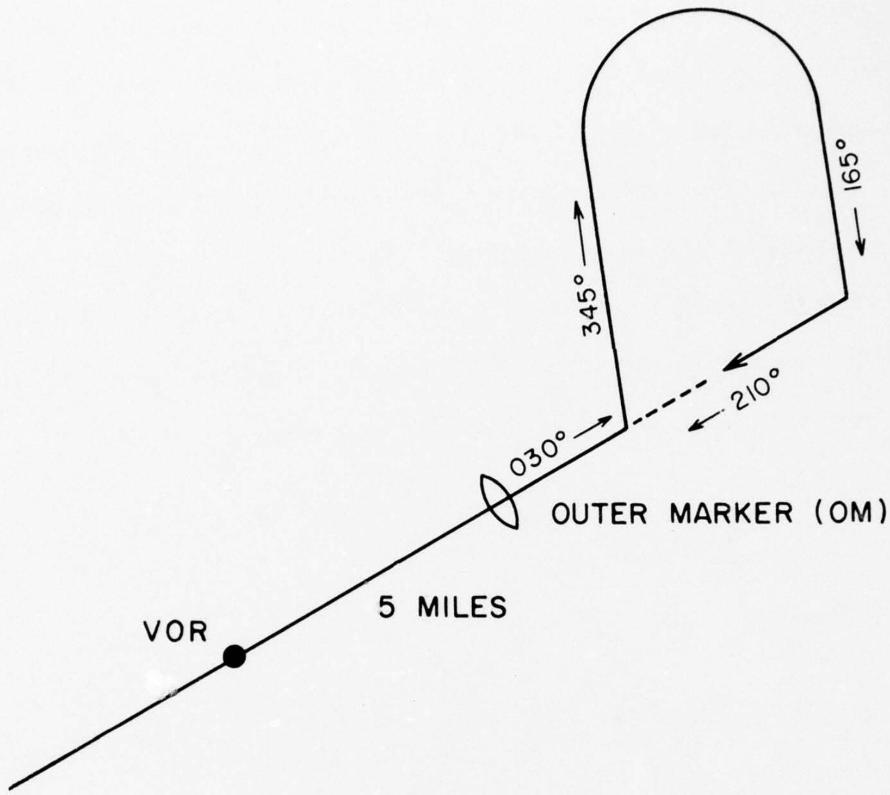

Figure 9. Typical procedure turn in an instrument approach.

Because there are no follow-up courses to use as a means of determining acceptability, all that is left is the use of a subjective questionnaire, asking experimental subjects to react to the PLATO instructional sequence.

In conclusion, this dissertation evaluates the PLATO program with the intention of deciding whether the approach has a potential that can be demonstrated in a realistic situation, and whether further research and evaluation is warranted. The evaluation is more formative in nature than summative, and wittingly does not address some important issues.

6. PLATO'S "INSTRUCTOR": AUTOMATIC EVALUATION AND FEEDBACK

One of the most appealing aspects of computers as a medium for teaching is the potential of using their computing power to evaluate how well a student is progressing and to provide him with meaningful feedback. Only if this is satisfactorily achieved will computers be regarded as teachers in their own right, rather than yet another teacher's aid.

In the traditional CAI lesson the criteria on which student performance is assessed are usually quite simple and readily available, generally taking the form of answers to multiple choice questions. In such cases the question plays the role of an error sensor (Crowder, 1960). Because there is usually a very small number of possible choices to any question, it is relatively easy to decide what to do with the student on the basis of his answer. (It is usually more difficult to decide what to do when the student has made a mistake than when he answers correctly). In some of the more advanced programs branching to remedial routines is based on the answers to more than one question, and remediation is often invoked only if the student has less than a certain percentage of all responses correct. Smallwood (1966) went a stage further and not only determined the path of each student through the program on the basis of immediate performance, but also used the performanced of all who had already taken the program in the branching strategy.

In a different context, Carbonell (1970) developed a CAI program based on a semantic net structure proposed and developed by Quillian, Collins and Warnock (see Quillian, 1968, 1969; Warnock and Collins, 1973).

In this program the student can make responses which are reasonable extrapolations of the previously given information and the system can "understand" them. Thus, although the computer does not have all possible answers stored in its memory, it can exhibit some "intelligence" by making inferences from its knowledge base.

The holding pattern program requires a radically different approach to student evaluation because the task is continuous and not discrete, both in terms of the time and the range of responses. The range of possible errors is greater in the sense that identical performances may result from one or more different performance deficiencies. For example, if the inbound leg were flown consistently to the left of course, possible reasons for this could be incorrect compensation for the crosswinds or perhaps incorrect use of the radio navigation equipment. Consequently automatic evaluation of this kind of performance is very difficult and no wholly satisfactory method has yet been found.

As was mentioned briefly in the previous chapter, the empirical validation of the automated evaluation of student performance proposed in Trollip (1975), revealed several serious deficiencies. Consequently another scheme was devised and successfully validated.

The original method involved establishing a "target area" at the end of the top turn. This was done for each pattern by calculating exactly where the end of the top turn would be if the pattern were flown perfectly. This point was then to be enclosed in an empirically established area designated as the "target area". If the pilot was within this "target" at the appropriate moment, then his performance

would be judged as being adequate, and if he failed to be in the "target" at the proper time, this would indicate that errors had been made in the execution of the pattern. Figure 10 illustrates a possible "target area" for a crosswind pattern.

In the validation of this method it turned out that failure to be in the "target area" certainly indicated that some errors had been made in the pattern. However, the inverse, that a plane hitting the "target" at the correct time indicated a good pattern, was found to be invalid. On many occasions, after being in the "target area" at the correct time, the student made errors later in the pattern which were not reflected in the evaluative judgement. Further, it was found that even though such deviations would have shown up on the pattern following, that inconsequential deviations from a well flown pattern often did too. Thus this method of evaluating student performance was not at all successful and was abandoned in favor of a new system.

The new method was more complex and more heavily dependent on observations and experience in the establishment of the criteria than was the initial approach. As the plane flew around the pattern, information was stored as to whether it was satisfying various "local" criteria. A "local" criterion is one that has only local significance, but when taken together with several others it contributes to an overall assessment of performance. Thus, one such "local" criterion was whether the plane was within half a mile of the VOR when the TO/FROM flag changed from TO to FROM at the end of the inbound leg. Alone this information contributes little to the overall evaluation of performance, but when placed into the context of all others it may provide important information.

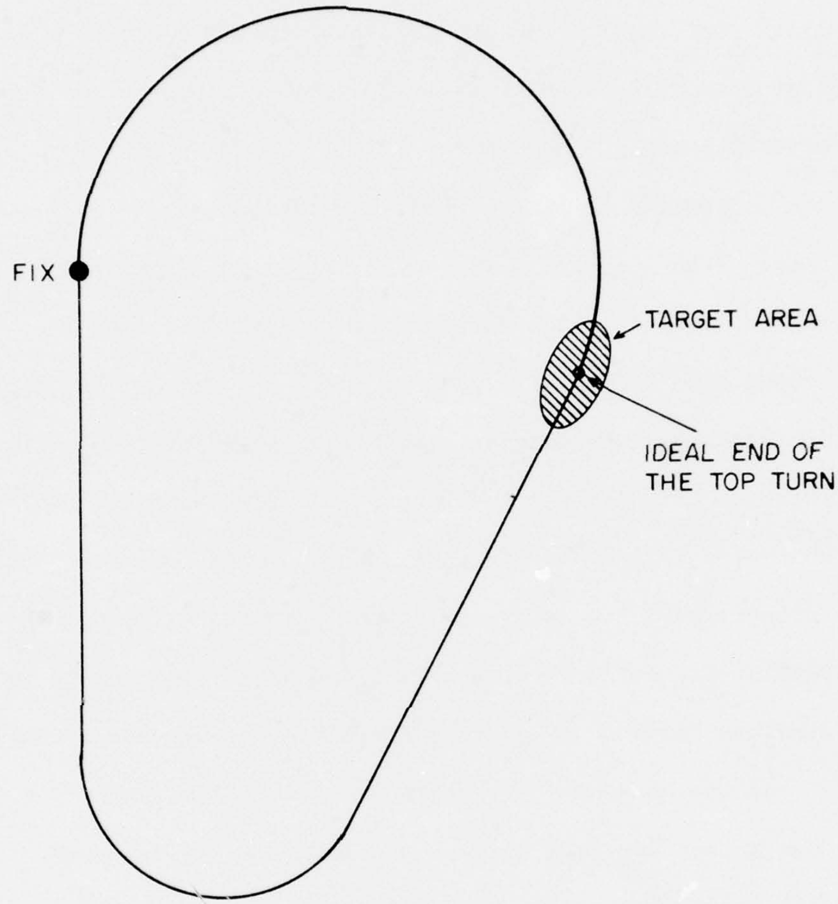

Figure 10. Illustration of the "target" method of evaluating performance.

For each pattern the following "local" criteria were applied:

1/ Whether or not the plane was within half a mile of the VOR when the flag flipped;

2/ Whether or not the top turn was flown at standard rate;

3/ The average heading on the outbound leg;

4/ Whether or not the plane was more than 5^o off the inbound radial when approximately three-quarters of a mile from the VOR inbound;

5/ The total time for the whole pattern.

The criteria mentioned in items 4 and 5 above provide more information about the overall performance than do the first three. However, the first three provide essential information about what may have caused violations of items 4 and 5. Thus if the plane were more than 10^o off to the right on inbound, and it had been more than half a mile to the right of the VOR on passing it, a good possibility would be that the initial lateral deviation caused the error on the final inbound, and not necessarily subsequent pattern errors. Consequently the overall assessment takes into account the contribution of these intial "flying" (as opposed to cognitive) errors.

On completing each pattern the student receives the following feedback as to his performance. First there is a factual reporting of what went wrong; for example, "Your rate of turn in the top turn was not good. You should have turned at 3^o per second, but you actually turned at 4.2^o per second which was too fast," or "When you were half a mile from the VOR you were more than 5^o off course to the right, Your inbound heading at this point was 010^o which would take you further from your course."

After this factual kind of feedback, the program attempts to provide pertinent information as to the cause of the various errors. The following example of how feedback is provided about bad timing in the no-wind pattern will illustrate the underlying approach to the evaluation.

The "local" timing criterion requires that the pattern be flown within twelve seconds from the ideal of 240 seconds. If the timing is satisfactory no subsequent feedback is provided on the timing aspect of the pattern. If the pattern is flown more than twelve seconds faster than the ideal, two causes are possible. First, the top turn could have been flown at a rate greater than standard. Second, if the rate of turn was adequate the likely reason is that the outbound leg had been flown too short.

On the other hand, if the pattern takes too long to fly, there are four possible causes. First, if the top turn was flown at less than standard rate, this itself is the likely reason for the bad timing. Second, if the top turn was flown at greater than standard rate <u>and</u> the plane was more than 10° to the right of the inbound radial approaching the VOR, the probable reason for the bad timing is that overcompensation for the bad top turn could have caused the plane to fly a distorted and lengthened course (Figure 11a illustrates this possibility.) A third situation in which the pattern takes too long to fly is one in which the top turn is flown at faster than standard rate <u>and</u> the plane ends up more than 10° off the inbound radial to the left <u>and</u> the plane's heading is taking it further off course. Here the pattern is distorted because the plane is flown across the inbound radial and no correction made (see Figure 11b).

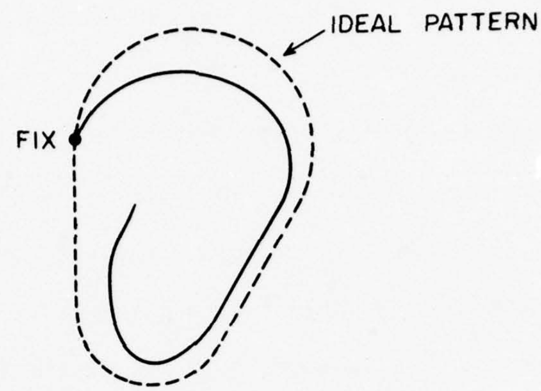

(a) Too much compensation for greater-than-standard-rate top turn.

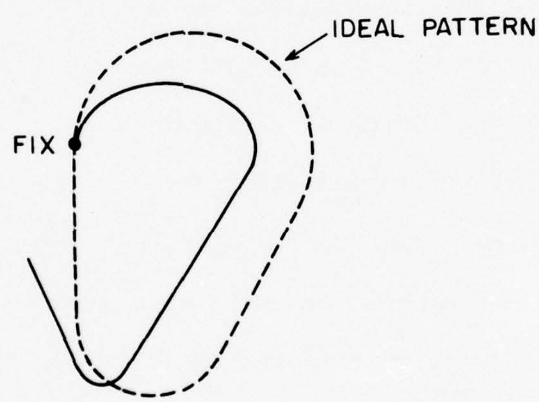

(b) Insufficient correction on inbound for greater-than-standard-rate-top turn.

Figure 11. Some possible errors in flying a holding pattern.

Finally, if none of the above situations has happened, the remaining situation resulting in the pattern being excessively long, is that the outbound leg was flown for too long.

While the previous discussion dealt with timing problems in the pattern, a similar approach is used in providing feedback in the case where the plane is more than 10° off course approaching the VOR on the inbound leg. Because there are more variables which can contribute to this error situation than in the previous example, a verbal description of all the possibilities is confusing. Consequently the various possibilities are illustrated in Figure 12.

Because this evalution model is based heavily on observational data, certain unexpected errors can not be diagnosed. In these rare cases the student is given a possible explanation, but warned that it may not be correct. In general, however, this quasi-Baysean approach is able to provide meaningful feedback under most circumstances.

The method of evaluation discussed in this chapter is but one way of attacking the very difficult task of providing feedback on complex, continuous performance. While it may not be the optimal method of assessing such performance, it represents an important advance in CAI because it has demonstrated that such feedback is both possible and useful.

Off course to the left on inbound: ("Approaching the VOR inbound, you were way left of course.")

Initially off course:			Rate in top turn			Outbound Heading			
Left	OK	Right	Fast	Standard	Slow	$>4^o$ left	OK	$>4^o$ right	Feedback
X				X			X		You flew the pattern well enough, but did not compensate for being to the left of the VOR initially.
X				X				X	Your outbound heading which brought you closer to the radial, as well as your initial position to the west of the VOR, resulted in you being to the left of course.
X			X				X		Your initial position west of the VOR was not helped by a greater than standard rate turn at the top.
X			X					X	A faster than standard rate turn at the top, and an outbound heading which would take you further to the west both just exaggerated your initial bad position.
X					X			X	Even though your slower than standard rate turn should have compensated for your initial bad position, you flew your outbound leg badly.
	X or	X		X				X	It seems that your outbound heading took you too far to the west resulting in you being to the left of course on the inbound leg.
	X or	X	X				X		It is possible that because your top turn was flown at greater than standard rate you landed up to the west of the inbound radial.

Figure 12. Feedback given under different circumstance for being more than 10^o off course on inbound.

Initially off course: Left	OK	Right	Rate in top turn Fast	Standard	Slow	Outbound Heading $>4^o$ left	OK	$>4^o$ right	Feedback
	X or	X	X					X	A faster than standard rate turn and a bad outbound heading took you to the west of inbound radial.
	X or	X			X	X			Your outbound heading must have been very bad to overcome a slow top turn and an initial position to the east of VOR.

Off course to the right on inbound: (Approaching the VOR inbound, you were way left of course.")

Initially off course: Left	OK	Right	Rate of top turn Fast	Standard	Slow	$>4^o$ left	OK	$>4^o$ right	Feedback
X					X		X		Although you were to the west of the VOR at the start of the pattern, it appears that a slower than standard rate turn in the top resulted in you being to the east of course.
X						X			Your outbound heading took you away from the inbound radial resulting in you being to the right of course on inbound.
		X		X		X			You started off the pattern to the east of the VOR, and then your outbound heading took you even further from it.
		X		X			X		You flew the pattern well enough but because you were to the east of the VOR to start with, you ended up to the east again.

Figure 12. Continued.

Initially off course: Left	OK	Right	Rate in top turn: Fast	Standard	Slow	Outbound Heading: $>4^o$ left	OK	$>4^o$ right	Feedback
		X		X		X			Although you seem to have corrected for your initial position to the east of the VOR, it does not appear to have been enough.
		X			X				You started the pattern to the east of the VOR, and so a slow rate of turn took you further away, and merely made the problem worse.

The catch-all feedback was:

It is difficult to say exactly what you did wrong. Try to keep a mental picture of what you are doing. This should help performance.

Figure 12. Continued.

7. METHOD

Four groups of private pilots were taught to criterion in a Singer Link General Aviation Trainer 2, called the GAT-2 (Figure 13), how to fly holding patterns in different wind conditions. The experimental conditions differed as follows: two groups (the control groups C1 and C2) were given prior groundschool instruction on holding pattern procedures by a lecturer in a classroom setting; the other two groups (the experimental groups E1 and E2) were taught in their groundschool by the computerized lesson on the PLATO CAI system. Groups C1 and E1 were then taught to fly holding patterns in the GAT-2 until they attained a prescribed level of competence, following which they were tested on a procedure turn to determine whether the strategies applicable to holding patterns transferred to this new task. Groups C2 and E2 were first tested on the procedure turn, following which they were taught to criterion on the holding patterns. Thus the experimental design may be illustrated as in Figure 14.

Subjects

Subjects were recruited for the study by advertising in the local student newspaper, The Daily Illini. Respondents were then sent a letter outlining the requirements of the study and asked to return a questionnaire if they could meet the requirements. Seventy-two questionnaires were returned. Two of the respondents withdrew before the study started because they had dropped out of school and were leaving the community.

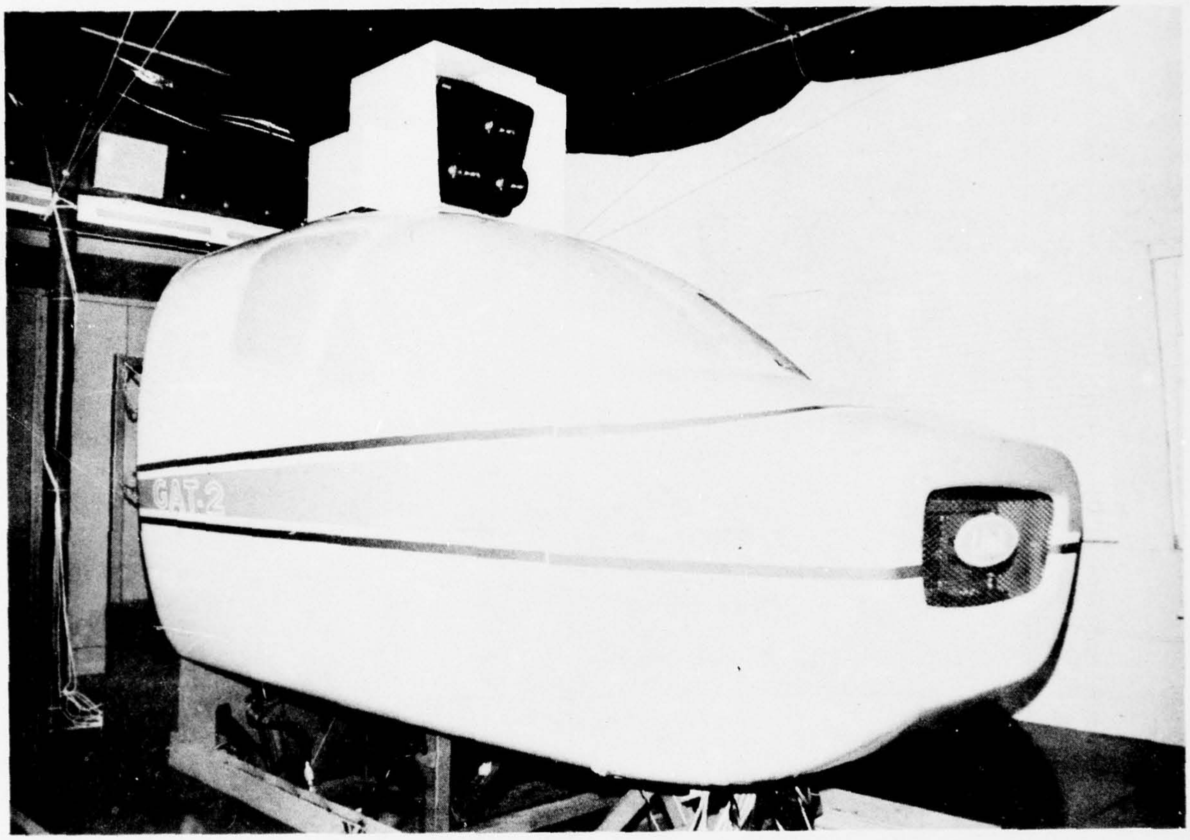

Figure 13. The Singer-Link General Aviation Trainer GAT-2.

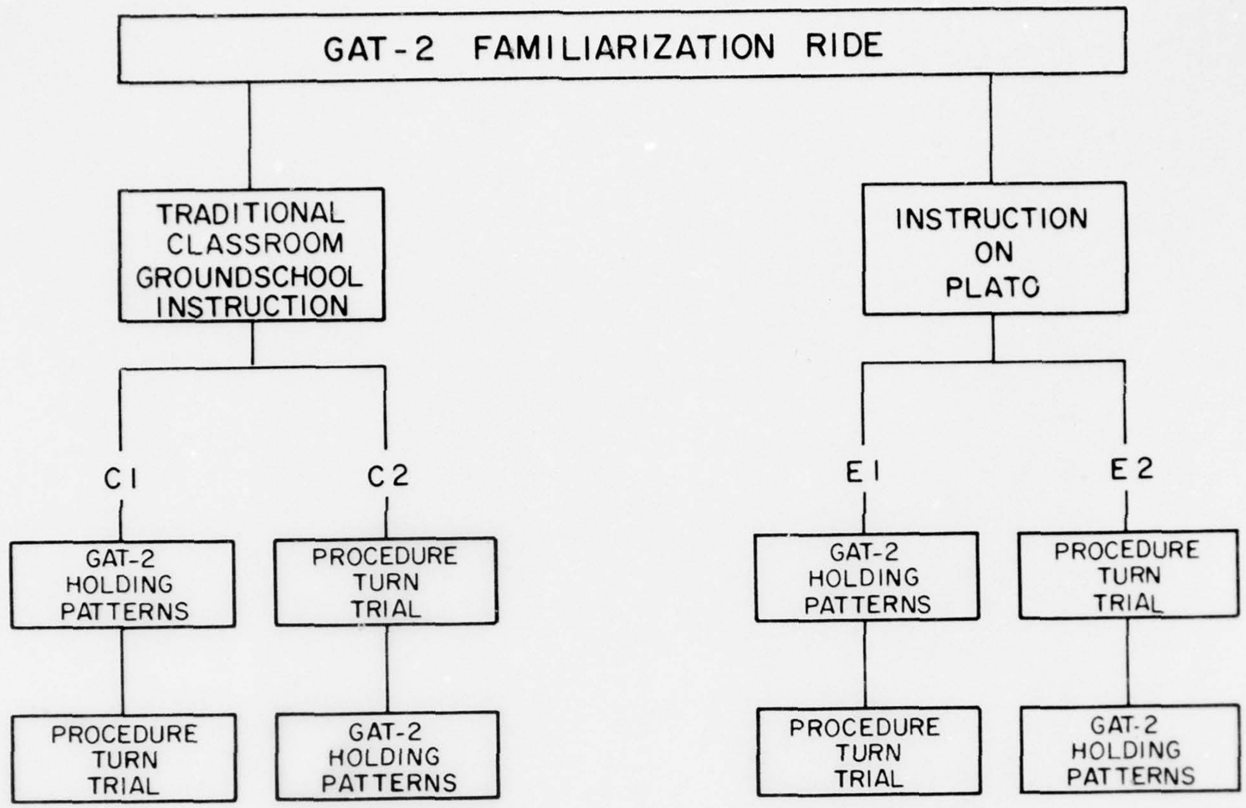

Figure 14. The overall study design.

All subjects were sent an information sheet briefly outlining the power settings of the GAT-2 for the various stages of flight. This was necessary as most subjects had only flown constant pitch propellor planes rather than the more complex constant speed propellor of the GAT-2. Each subject then completed a familiarization ride in the GAT-2 which required him to make an instrument take-off, complete various basic flight maneuvers, and fly a tracking task using the same radio navigation facility (the Very high frequency Omni-Range, or VOR as it is called) to be used in the final study. With the exception of the instrument take-off, criterion had to be attained on each task before progress was permitted to the next.

The first of these basic flight maneuvers was to fly the GAT-2 straight and level for one minute without deviating more than 10° from the desired heading, without exceeding 10° of bank, and without being more than 150 feet off the prescribed altitude. The second task was to perform standard rate turns in both directions to specific headings. A standard rate turn is defined as a turn executed at 3° per second, and is indicated in the cockpit by a needle and appropriate index marks. In this task the rate of turn had to be held to within one degree per second, the altitude to within 150 feet, and the roll out heading to within 10° of the desired heading. The final task required the subject to fly inbound to a VOR on a given radial, and on crossing the VOR to reset the omni-bearing selector (OBS) on the navigation receiver to another radial, to intercept this new radial within three miles of the VOR and to fly outbound on it. As before the permissible deviation was 150 feet in altitude. The plane also had to stay within 10° of the inbound radial.

On each of these tasks, if a violation of the criteria occurred the student was allowed to complete the assignment, and was then debriefed by the instructor. The GAT-2 was reset to its initial position, heading and altitude. Progress from one task to the next was accomplished by performing two successive errorless trials.

The GAT-2 is sensitive to very small changes in pitch, and inexperienced pilots are sometimes hard-pressed to maintain altitude. The familiarization flight not only ensured that all subjects were competent to fly the GAT-2 within reasonable limits, but also provided measures to be used in the data analysis as covariates to reduce variance attributable to individual differences.

Grouping

From the original group of seventy respondents eighteen were selected to form a trial group whose purpose was to allow the experimenters to perform a complete dress-rehearsal of the entire study. Several "bugs" that had not appeared in the previous pretesting were found and eliminated. This dress-rehearsal also gave the necessary practice to the instructors and observers to standardize their techniques.

The remaining fifty-two subjects were rank-ordered on the basis of their performance on the familiarization flight, and split into two groups at the median. From each group thirteen subjects were randomly assigned to be control subjects and thirteen experimental subjects. The twenty-six control subjects were then randomly divided into two groups C1 and C2. The two experiemtnal groups E1 and E2 were formed in like fashion.

Fifty subjects finally completed thier assignments. Both of those not finishing (due to personal time constraints) were in Group E1.

Observers

The four observers used in the study were certified pilots. Three had instrument instructor certificates and took care of all instruction; the fourth was a private pilot with extensive instrument flying experience. All observers were trained carefully before the main study as to the procedures to be used and the manner in which the booklets should be marked. The observer in the front right seat of the GAT-2 (see Figure 15) acted as the flight instructor. He was responsible for giving the subject appropriate instructions and for stabilizing the GAT-2 before each trial. The observer in the back seat was responsible for resetting the GAT-2 before each trial to the appropriate position, altitude and heading. He was also responsible for changing the strength and direction of the wind between tasks.

Ground School Instruction

All control and experimental groups received ground school instruction on holding patterns before being taught in the GAT-2. The control groups were given a lecture on holding patterns by a member of the Institute of Aviation's Pilot Training Department, who was an Instrument Flight Instructor with extensive experience in conducting ground schools. This lecture, delivered in the standard manner, lasted seventy-five minutes including questions from the students.

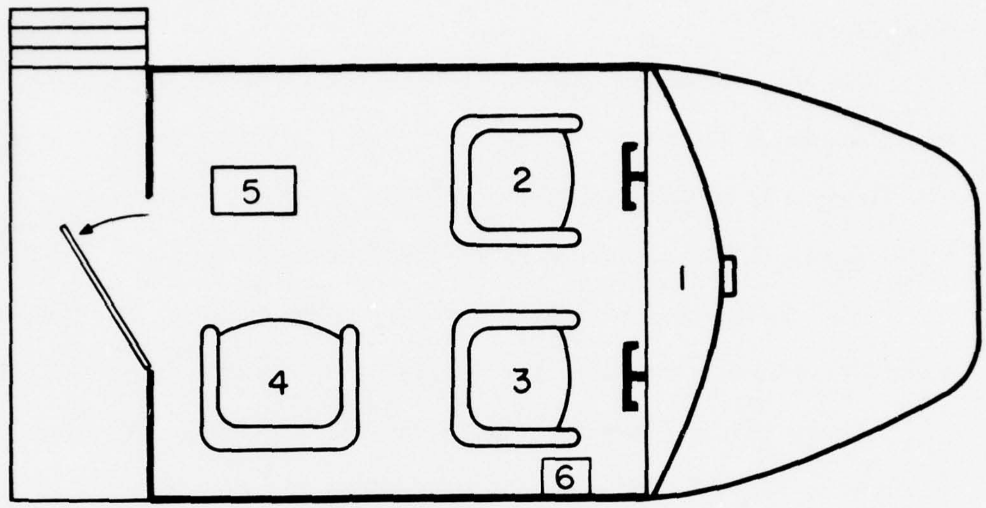

1. Instrument Panel
2. Subject's Seat
3. Instructor's Seat
4. Observer's Seat
5. Control Panel
6. Distance Measuring Equipment (DME) out of Subject's Sight

Figure 15. Schematic of interior of GAT-2.

The experimental groups were given equivalent ground school instruction by the PLATO program, and also required to fly several holding patterns using PLATO's simulated plane and the interfaced hand-controller discussed previously. Subjects used the PLATO lesson individually and, except for the experimenter being available to handle any unforeseen system malfunctions, were unsupervised.

The GAT-2 Sequence of Tasks

Each subject was required to complete a training session in the GAT-2 on holding patterns in which progress through the session was based on the student achieving criterion at each stage. Each subject also listened to tape-recorded instructions on how to fly procedure turns in an instrument approach, following which he had to fly one approach in strong wind conditions. Groups C1 and E1 flew the procedure turn before the holding pattern sequence, while the groups C2 and E2 had the order reversed.

The Holding Patterns

The holding pattern training sequence consisted of four parts. First, the subject was required to fly a series of holding patterns without any instruction in a difficult wind condition. Second, he was given instruction to fly holding patterns in no-wind situation, followed by a left crosswind condition, and finally with a strong right quartering tailwind. Figure 16 illustrates the structure of this sequence.

The purpose of the first task in which the subject received no prior information about wind conditions other than that there was a wind, was to evaluate immediately how much had been learned from the ground school instruction. While the subject was flying the instructor was not allowed

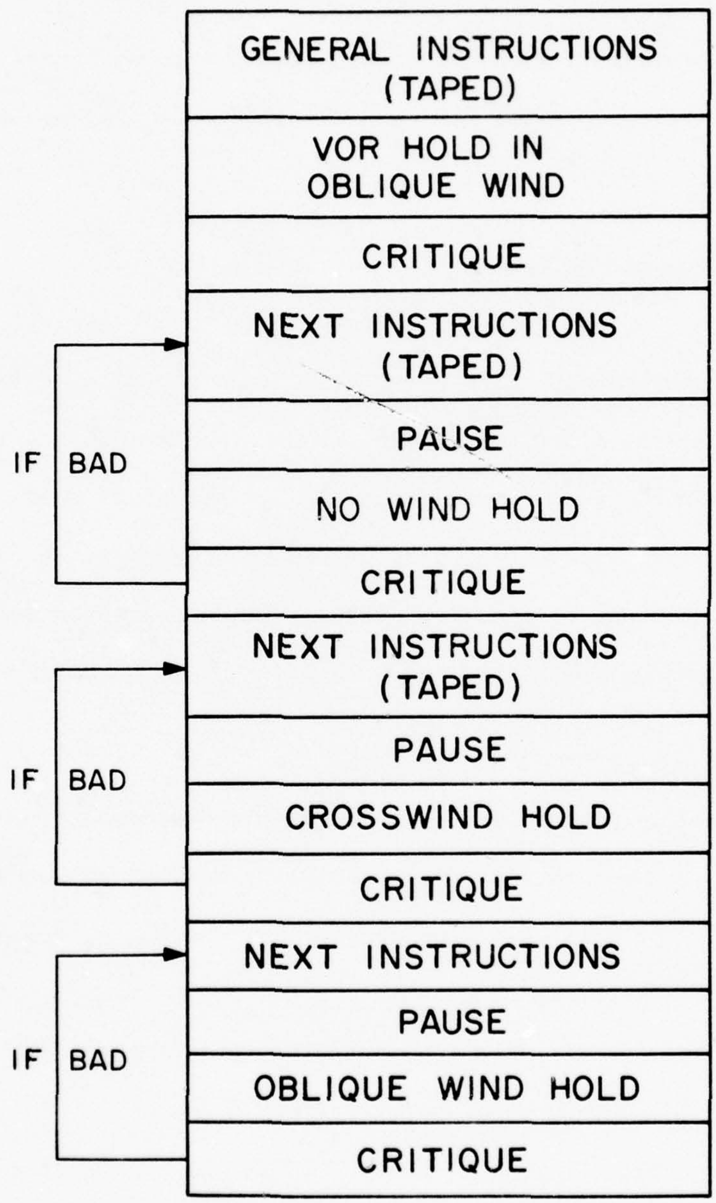

Figure 16. Structure of the holding pattern GAT-2 training package.

to make any comments or criticisms. In this task the subject flew either until he had completed three patterns or until he indicated he was lost. In a few cases the instructor took over control when, in his opinion, after a sufficient amount of time the subject was completely disoriented but had not yet voiced that opinion.

At the end of this task, whether terminated due to being lost or because three patterns had been completed, the instructor debriefed the subject completely. All mistakes were pointed out both in terms of aircraft control and pattern strategies.

The wind in this "test" pattern was a strong quartering tailwind from the holding side. This is generally recognized to be the most difficult condition in which to fly holding patterns because failure to compensate markedly shortens the inbound leg and at the same time blows the plane well off course.

The final three stages constituted a training sequence starting from a no-wind condition to the same right quartering tailwind situation used in the "test" pattern, except that the wind became even stronger. In each case the subject was given taped instructions on how to fly each pattern. He was told what wind to expect, and was fully debriefed at the end of each trial by the instructor. As before, the instructor was restrained from making any comments while the patterns were being flown. Progress from one wind condition to the next was achieved by flying two successive patterns within the criteria outlined in the following paragraphs. Successful completion of the final task was taken, for purposes of this study, to indicate that the subject could fly holding

patterns competently. It must be mentioned, however, that other aspects of flying holding patterns in the real world, such as entry procedures, and holding at intersections of VOR radials, were not part of this study nor of its training program.

The Criteria

There are two distinct types of errors related to the planning and execution of holding patterns. The first type of error results from the incorrect planning of pattern strategies, and is manifested in deviations from the ideal pattern shape. The second type of error results from an inability to choose the strategies, whether right or wrong, while flying the plane at the same time. In this case control of the plane deteriorates, especially with respect to altitude. In the ideal case, the pilot would be expected to choose and apply the correct strategies while maintaining full control over the plane.

In fact there is a third type of error that could appear. This is the kind of error resulting from an inability to fly the plane accurately irrespective of what else is happening. This type of error varies from individual to individual and makes the analysis of interesting errors more difficult. Consequently it is important to reduce the likelihood of these sorts of errors from appearing.

One method to do this is to ensure that all subjects have the ability to fly within certain tolerances. This was done in the GAT-2 Familiarization Flight. In all maneuvers, for instance, altitude had to be held within 150 feet from the given altitude. A second way of reducing the intrusion of control errors into the data is to select criteria that were

sufficiently lenient that ordinary control errors would not violate them, but sufficiently sensitive that the addition of other errors would be noted.

In extensive pretesting it was decided that the altitude criterion should be 200 feet from the given altitude. From the Familiarization Flight it was known that all subjects could control the GAT-2 within 150 feet and that if the pilot's attention were diverted that the ensuing error was usually large. The 200 foot criterion allowed some irregularities but also was violated when the pilot let his attention wander from the task of flying to the task of working out what to do.

From the performance of experienced pilots on the holding pattern tasks it was decided that a ten second tolerance could be allowed for the time of the inbound leg, and that when about a mile from the fix on inbound the plane should not be more than ten degrees off course.

Once these tolerances had been decided upon, it was necessary to incorporate them into a scoring scheme that could be easily used by the observers, that was not unfair on the subjects in terms of penalizing them unnecessarily for minor errors, and that provided the necessary information to the experimenter as to the performance of the subjects on the holding pattern task as opposed to their control of the GAT-2.

After considerable pretesting the following scheme was implemented which seemed to satisfy all the requirements. First, the pattern was broken down into five parts: the approach to the fix before entering the hold; the top turn; the outbound leg; the bottom turn; and the inbound leg. All parts were scored with respect to altitude into three

categories: no mark was made if the altitude remained within 100 feet of the given altitude; if the GAT-2 varied more than 100 feet from altitude a box was checked. If the altitude varied by more than 150 feet a second box was checked, and finally if there was more than 200 feet deviation, a third box was checked. Figure 17 shows how the boxes were marked in different cases. Note how observers distinguished between deviations up or down.

Deviations from the desired flight path on a horizontal basis were scored only on the initial approach and the inbound leg. When the GAT-2 was three-quarters of a mile from the holding fix, as indicated on Distance Measuring Equipment in the GAT-2 which was concealed from the subjects sight, the position of the course deviation indicator (CDI) on the VOR was noted. If the needle was on the colored bands, no marks were made; if lying between the bands and the pegs on the edge the box indicating greater than five degrees off course was marked (with left or right indications as well); and if the needle was pegged the box was marked indicating more than ten degrees off course. Actually the pegged position was a little more than ten degrees off course; about twelve to be exact. Figure 18 illustrates the course marking system.

Finally if the inbound leg were flown such that it was more than ten seconds longer or shorter than a minute, this was taken to be a violation.

All these deviations were classified into two types: errors and critical errors. Critical errors were more than 200 feet off altitude, more than ten degrees off course and more than ten seconds off on the

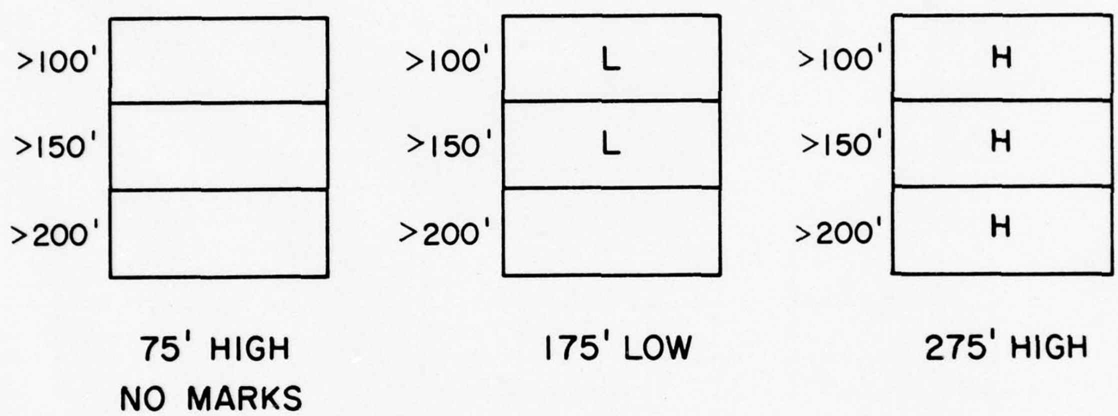

Figure 17. Scoring altitude deviations.

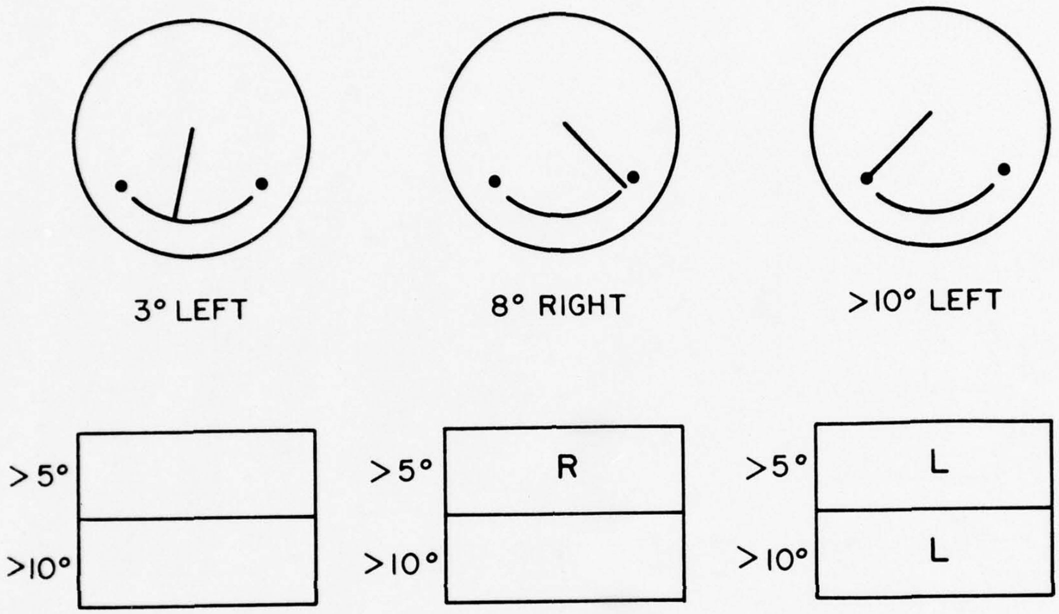

Figure 18. Scoring deviations from inbound radial.

minute inbound leg. Errors were scored purely for informational value while critical errors were used to determine whether criterion had been attained on any trial.

Some further modifications were made for logistical reasons. Because it was unfair to penalize a student twice who was off on his altitude at the end of the top turn, say, who started to correct, but was still off on his altitude when reaching the outbound leg, a short "grace" period was built in at the beginning of the leg. Similar periods were built into the beginning of the top turn and the beginning of the bottom turn too. Because the inbound leg is the most important leg, no leeway was allowed on it. The "grace" period in the turns was ninety degrees of turn, and on the outbound leg fifteen seconds after rolling the wings level.

To achieve criterion on each holding pattern then, the subject had to fly the top turn, outbound leg, and bottom turn with no more than a total of one critical error, and the inbound leg with no critical errors. In the event of the two observers disagreeing as to whether criterion had been attained, it was assumed that it had not, and the trial regarded as one with errors. The original scoring of such patterns remained. The approach leg was not used as a measure of criterion performance. Figure 19 details the criteria for the pattern. The criteria remained the same irrespective of the wind condition.

In retrospect, this scheme worked out well. It was sufficiently sensitive to catch those errors that were to provide data, and sufficiently lenient that the poorer pilots could remain within the criteria if they knew what they were doing.

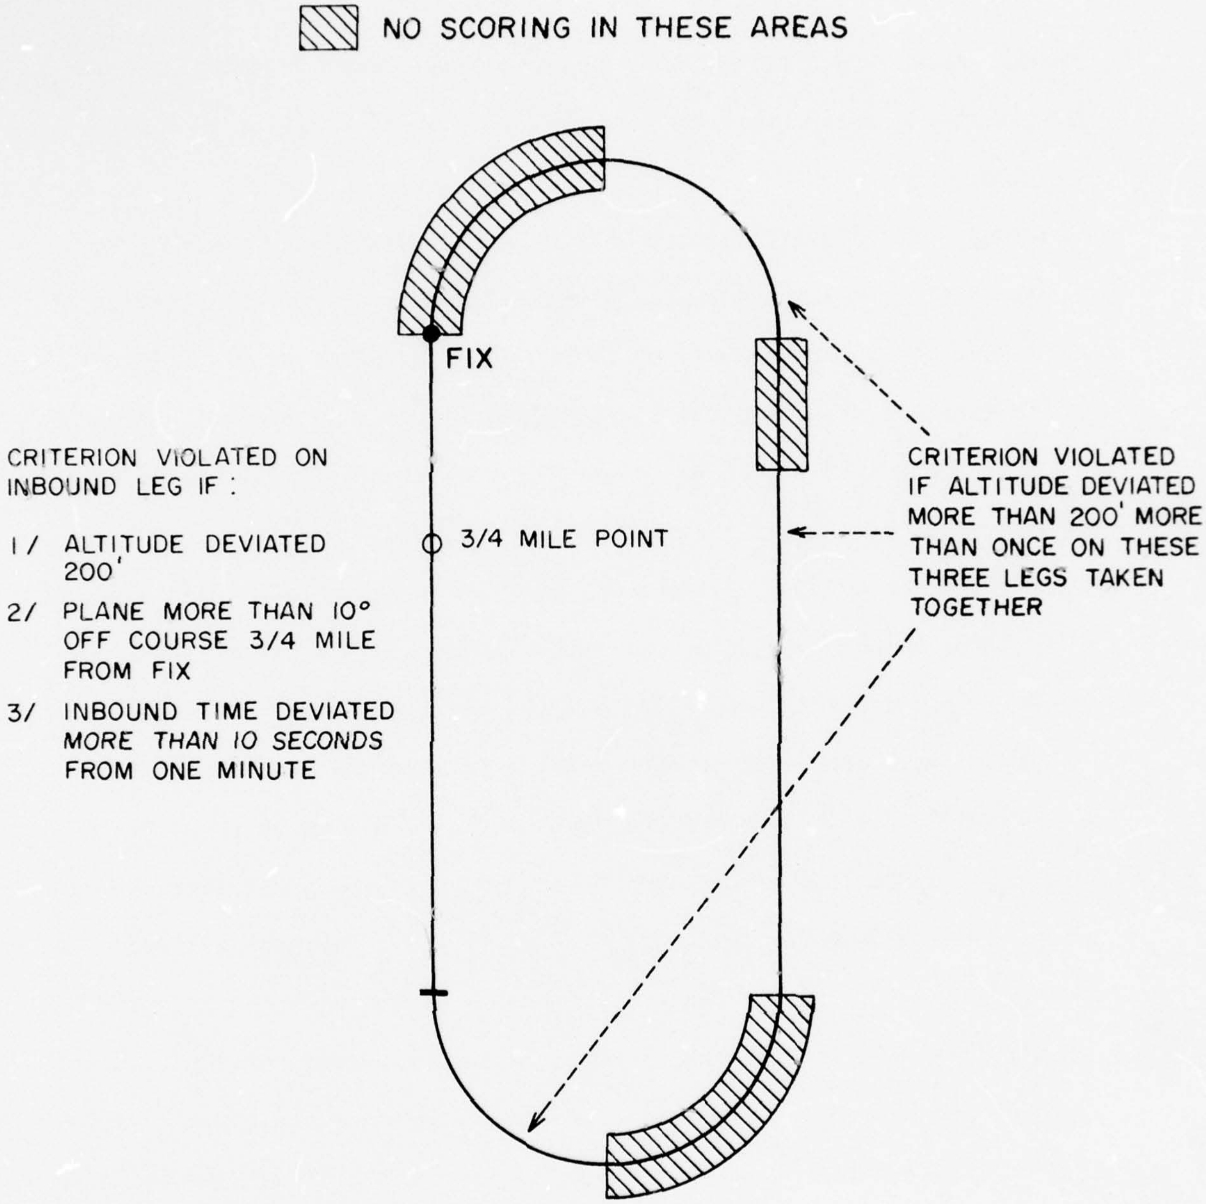

Figure 19. Criteria for the holding pattern.

The Procedure Turn

As discussed in Chapter 5 the inclusion of the procedure turn in the study was to provide a means of measuring the transfer of the skills acquired in learning holding patterns to a task similar in nature, but different in execution. Because the procedure turn, like the holding pattern, is dependent on a ground fix and requires a reversal of direction, it too demands compensations to be made for winds. Without such compensation the plane may be blown so far off course that reasonable recovery is impossible, a situation which would necessitate another attempt at best, or cause disorientation and panic at worst.

Each subject was played tape-recorded instructions explaining all the necessary details for flying the procedure turn. He was encouraged to spend some time thinking about how various winds would affect the execution of such a turn, and what strategies would have to be invoked to compensate for the resulting drift. When the subject indicated that he had thought sufficiently about the task, he was then told the surface winds in the vicinity of the procedure turn as a hint of the winds aloft, and asked to fly one complete procedure turn. (See Figure 20).

Similar scoring measures were used in the procedure turn as were used in the holding pattern. The subject was given only one attempt at the turn without any feedback during performance from the instructor. The purpose was to be established on the inbound radial before crossing the outer marker inbound (See Figure 9).

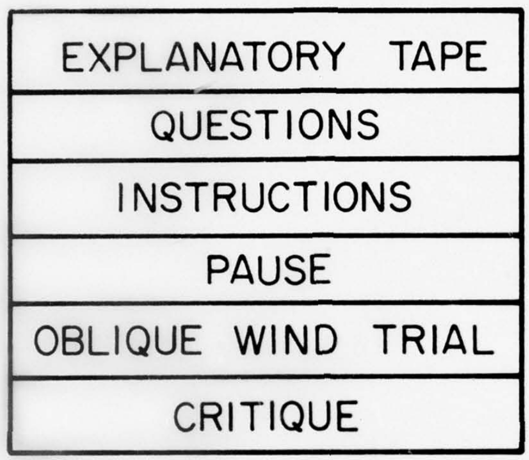

Figure 20. Structure of the GAT-2 procedure turn test.

Procedures

Each subject was required to participate in four sessions. The first was the familiarization ride, the second the ground school instruction (either a lecture or the PLATO sequence),the third either Holding Patterns I which consisted of the "test" pattern and the no-wind pattern, followed by Holding Patterns II containing the crosswind and quartering tailwind sessions, and the Procedure Turn (Groups C1 and E1), or Procedure Turn and Holding Patterns I, followed by Holding Patterns II (Groups C2 and E2). The subjects received no remuneration for their participation, but their flight logs were endorsed appropriately by the flight instructor in charge of each session.

8. RESULTS

Analysis

For the most part the analysis of the data was accomplished by using a two-factor analysis of variance design. The first factor was the method by which the subjects received their ground school instruction; either the traditional lecture or by PLATO. These will be called the treatments. The second factor was the order of presentation of the holding pattern sequence and the procedure turn. This will be called order. Figure 21 illustrates this design showing how the previously discussed groups (C1, C2, E1, E2) are distributed in the design.

In a few cases an anlaysis of variance is performed on dichotomous data such as whether a subject was lost or not. Doing this is a point of considerable controversy in statistics and the F-ratios generated by this method are open to some doubt. Justification for this approach stems from the implication inherent in Cochran's (1950) indication that the F statistic computed by treating the data as if they were normally distributed, will yield probability statements close to those obtained by use of the dichotomous Q statistic (Winer, 1975, p. 305).

Most of the data analysis was performed on the University of Illinois IBM 360 computer using the SOUPAC (1974) statistical package, or by the PLATO statistical program.

Homogeneity of Groups

Analysis of variance on the number of errors made on the familiarization flight (the basis used for grouping) indicated that there were no significant differences between the groups. There were also no significant

		TREATMENT	
		TRADITIONAL	PLATO
ORDER	HOLDING PATTERN PROCEDURE TURN	GROUP C1	GROUP E1
	PROCEDURE TURN HOLDING PATTERN	GROUP C2	GROUP E2

Figure 21. Analysis of variance design.

differences between the groups on the basis of the number of hours flown before the study. Table 1 presents the composition and distribution of the groups by errors on the familiarization flight, and by hours flown.

Student Performance on the PLATO Program

As was mentioned before, when the subject arrived at the PLATO terminal he first read a handout before signing onto the system. After signing on, he then completed the first part of the program which quizzed him on what he had read; then he was put through a training sequence on holding patterns, at each stage being required to acheive PLATO's criterion. For Group E1 the mean time taken from beginning to end, including the reading of the handout was 111.1 minutes. For Group E2 the time was 101.6 minutes. The difference between the groups is not significant ($p > .10$).

Unfortunately due to a programming error, some of the data collected while subjects were on the PLATO system were lost, including the time spent on each of the pattern conditions, and the total time spend "flying" PLATO's plane. However, correlation coefficients between total time spent, as reported in the previous paragraph, and various variables of performance in the GAT-2, indicate that performance in the GAT-2 correlates quite highly with performance in PLATO.

In general, the highest correlations occurred with those maneuvers in the GAT that required more than the very basic instrument skills. Thus, for instance, the correlations between "time on PLATO" and "errors made on the instrument take-off" (a very complex instrument maneuver) for the two experimental groups are .497 and .471. For the number of trials

TABLE 1

Composition of Groups by Score on Familiarization Flight and by Flight Hours

1 = Control C1   3 = Experimental E1

2 = Control C2   4 = Experimental E2

												4															
												4															
		4								4		4															
		4								2		3				2											
		4				3		3		2		3		4		1	4										
		3				3		2	3	1		2	2	2		1	3	4	2		3*			4			
3		1		2	1	2		2	1	1		1	1	1	3	1	2	3*	2		1		4	1	3		4
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27

Distribution of Groups by Errors on Familiarization Flight

Familiarization Scores				Flight Hours			
C1	C2	E1	E2	C1	C2	E1	E2
2	4	0	2	61	100	50	75
5	6	2	2	130	150	51	165
9	8	6	2	140	68	115	63
10	8	6	10	102	53	81	81
10	10	8	12	79	73	73	88
12	10	9	12	61	100	58	100
13	12	12	12	69	80	52	85
14	13	12	14	85	77	110	62
16	14	15	17	50	82	72	55
16	16	17	18	62	53	166	73
16	17	18*	23	80	72	90*	59
21	19	21*	24	83	94	130*	77
24	19	25	27	55	120	67	130

*(Subjects who failed to finish)

needed to attain criterion on the VOR tracking task in the familiarization flight the coefficients of correlation are .434 and .433, and for the crosswind pattern .365 and .422 (see Table 2). In all but one case the correlations between the time spent on PLATO and the performance variables are positive or very close to zero. In one case ("course errors" in the crosswind pattern) group E1 has a correlation of -.382 while E2 has one of .356. No reason can be found for this discrepancy.

The overall positive correlation between time spent on PLATO and performance in the GAT-2 suggest that similar skills are required in both tasks. This should be the case if PLATO simulates faithfully the appropriate cues used to fly under instrument conditions.

Student Reaction to the PLATO Program

With three exceptions the subjects reacted favorably to very favorably towards the PLATO instruction they received, especially the part in which they had to "fly" PLATO's plane. All three of the dissenting subjects commented that they had some problem becoming comfortable with the control characteristics of the hand-controller (five others also commented on this).

Nine of the subjects spontaneously commented that the PLATO lesson helped them visualize what they were doing, and thus helped their performance. Seven commented that the PLATO program gave them the opportunity to experiment with different strategies, and eleven commented on the fact that they could concentrate on what they were doing without being distracted by other things.

TABLE 2

Correlation Coefficients between "Time on PLATO" and Various Performance Variables for each of the Experimental Groups

Task	Variable	Group E1	Group E2
Familiarization flight	Take-off errors	.497	.471
	Trials to criterion (Straight and level)	.129	-.055
	Trials to criterion (Standard rate turns)	.305	-.150
	Trials to criterion (VOR tracking)	.434	.433
"Test" pattern	Number of patterns complete	-.177	.074
No-Wind pattern	Trials to criterion	.089	.206
	Altitude errors	.243	.332
	Course errors	.067	.505
	Timing errors	.295	.174
Crosswind pattern	Trials to criterion	.365	.422
	Altitude errors	.490	.366
	Course errors	-.382	.356
	Timing errors	.055	.361
Quartering wind pattern	Trials to criterion	.703	.232
	Altitude errors	.722	.016
	Course errors	.228	.293
	Timing errors	.059	.423

Reliability of the Performance Measures

From the scoring booklets data were summarized as follows: altitude errors (not including critical errors); critical altitude errors; course errors (not including critical errors); critical course errors; altitude errors on the inbound leg (excluding critical errors); critical altitude errors on the inbound leg; timing errors. So, if a subject flew six patterns before achieving criterion, his score in these various categories would be the sum of the errors made over all six patterns. This procedure was followed for both observers.

Also recorded for each task were the number of trials (patterns) required to attain criterion, and the time it took to do so. Similar data were tabulated for the approach segment of each pattern, that is for the flight between the initial starting point six miles from the VOR to the VOR. Thus for each of the three wind conditions, the raw data were reduced to twelve composite data points.

The "test" holding patterns were handled in the same way, except that "the number of trials to criterion" was replaced by "the number of patterns completed before the instructor took control." An indication was also added as to whether the subject had gotten lost.

The procedure turn task data were handled in similar fashion with approach and turn segments. Most important on this task was the indication as to whether the subject had established himself on the inbound radial before reaching the outer marker inbound.

The reliability of the performance measures was assessed by computing the Pearson product-moment correlations between the summary scores of each

of the observers. Tables 3-6 show what the data points were that were used in the analysis, and the observer-observer correlation coefficients of each of these data points. It should be remembered that these correlation coefficients indicate not how well the two observers agreed on any single error, but rather how their overall assessments agreed as scored over several patterns. All these coefficients were based upon fifty pairs of composite observations.

Because the reliability of the scores is high, the mean of the two observers' scores is used in the analysis. In cases like the number of "trials to criterion" or "time to criterion" the data recorded by the instructor is used.

Treatment Effects on the Performance of the "Test" Holding Pattern

The inclusion of this "test" pattern was designed to give an immediate indication, in the most obvious way, of whether there were any differences in treatment effects. The two measures providing this information are whether the subject got lost, and how many patterns he flew until lost. It should be noted that the trial was terminated after three patterns if the subject had not indicated disorientation. The number of subjects losing themselves, and the mean number of patterns flown before termination are shown in Figures 22 and 23. There was no significant difference in the numbers of subjects losing their position while the difference between treatments on the number of patterns flown before termination barely missed arbitrary statistical significance, $p = .056$ (Table 7).

Fewer members of the two experimental groups asked their instructors questions before taking control of the simulator than did their control

TABLE 3

Summary Data Sheet on Familiarization Flight with

Observer-Observer Correlation Coefficients

Take-off errors	.787
Number of trials on straight-and level task	1.000
Number of errors	.983
Number of trials on standard rate turn task	1.000
Number of errors	.962
Number of trials on VOR tracking task	1.000
Number of errors	.939
Number of altitude errors	.928

TABLE 4

Summary Data Sheet on "Test" Pattern with

Observer-Observer Correlation Coefficients

Seconds thinking		Not applicable*
Minutes to completion		Not applicable*
Lost		1.000
Number of patterns completed		1.000
Approach:	Altitude errors	.929
	Critical altitude errors	.814
	Course errors	.726
	Critical course errors	.750
Patterns:	Altitude errors	.951
	Critical altitude errors	.888
	Course errors	.941
	Critical course errors	.832
	Timing errors	.874

*These observations were made by the instructor only.

TABLE 5

Summary Data Sheet for Holding Pattern Training Sequence with Observer-Observer Correlation Coefficients

Minutes to completion		Not applicable*		
		Pattern 1	Pattern 2	Pattern 3
Number of trials to criterion		1.000	1.000	1.000
Approach:	Altitude errors	.978	.924	.851
	Critical altitude errors	.872	.932	.518
	Course errors	.861	.816	.868
	Critical course errors	.744	.808	.565
Patterns:	Altitude errors	.976	.941	.942
	Altitude errors on inbound	.937	.889	.914
	Critical altitude errors	.945	.923	.963
	Critical altitude errors on inbound	.956	.940	.970
	Course errors	.930	.936	.911
	Critical course errors	.953	.931	.924
	Timing errors	.930	.913	.924

*This observation was made by the instructor only.

TABLE 6

Summary Data Sheet for Procedure Turn Task with

Observer-Observer Correlation Coefficients

Time thinking	Not applicable*
Total altitude errors	.845
Total course errors	.661
Crab established inbound	1.000
Altitude errors	.844
Critical altitude errors	.906
Established on inbound before OM	1.000

*This observation was made by the instructor only.

counterparts (Figure 24). However, the experimental groups thought longer about what they were to do than the control groups (Figure 25). Neither of these differences was significant.

Treatment Effects on the Performance of the No-Wind Pattern

No significant differences were found on any of the dependent variables between the control and experimental groups, neither as a factor of treatment nor by virtue of the order in which the holding pattern sequence and the procedure turn were presented. In all cases however, the mean number of errors made by the control groups was greater than that of the experimental groups. Also those groups who took the holding pattern sequence before the procedure turn scored worse than their counterparts. Figures 26-28 illustrate these trends for various variables.

Figure 29 shows the differences between groups for the total number of errors and the total number of critical errors to criterion on this no-wind task. Tables 8 and 9 show that there were no significant differences either by treatment or order.

Treatment Effects on the Performance of the Crosswind Pattern

The same trends discernable in the no-wind pattern showed in this crosswind pattern sequence. In all cases the mean number of errors made by the PLATO groups was less than that made by the control groups. The noticeable (but statistically insignificant) effects in the no-wind patterns of the order of presentation of the holding patterns and procedure turns weakened in the crosswind situation.

The PLATO groups took significantly ($p < .05$) less time than the control groups to attain criterion and made significantly fewer errors in doing so. There were marked differences in the number of trials to criterion and the number of critical errors although these differences were not significant.

Figures 30 and 31 show time and trials to criterion, and total errors and total critical errors to criterion. Tables 10-13 give the analysis of variance tables for these variables.

In addition to the above there was a very strong, almost significant, ($p = .058$) effect of order on critical altitude errors. This trend, however, was not in the predicted direction because those groups who had more practice (those which had flown the procedure turn before the holding patterns) made more errors. Figure 32 and Table 14 illustrate this.

Treatment Effects on the Performance of the Quartering Tailwind Pattern

All four groups performed almost identically on this task. As can be seen from Figures 33 and 34 treatment and order effects were very small and certainly not significant. Tables 15-18 give the analysis of variance for time and trials to criterion, and for errors and critical errors to criterion. The apparent interaction of treatment with order in time to criterion was not significant.

Treatment Effects on Performance of Entire Holding Pattern Training Sequence

Figures 35 and 36 illustrate the performance of the groups in terms of time and trials to criterion, and errors and critical errors

to criterion. Once again the experimental groups performed better than the control groups, although with the exception of critical errors, not significantly. The differences between groups in the case of critical errors were significant (p = .015).

There were no significant differences between groups on the basis of the order of presentation of the tasks. Tables 19-22 give the analysis of variance for the measures mentioned above.

Treatment Effects on Performance of the Procedure Turn Task

The overall differences on the procedure task came as a function of the order in which the procedure turn and holding pattern tasks were performed. In general the groups that had completed the holding pattern sequence before flying the procedure turn performed better on the turn than those who flew it as the first task.

The planning of the turn was measured by two major variables; whether the subject established a crab to counter the effects of wind prior to reaching the OM outbound; and, more important, whether he was established on the inbound radial before reaching the OM inbound. The effect of order on the establishment of the crab barely missed the .05 level of statistical significance (p. = .053), and on being established on inbound was significant at the same level (p = .011). Figures 37 and 38 illustrate these differences. Tables 23 and 24 give the analysis of variance tables. In both cases the effect of treatment was not significant.

Control of the GAT-2 during the procedure turn was measured by using the same altitude limits as in the holding pattern sequence. There were very reliable differences between groups by order of presentation ($p < .001$) on both the number of altitude errors and the number of critical altitude errors made during the turn (Figure 39 and Tables 25 and 26).

Treatment Effects on the Performance of the Holding Pattern Training Sequence on a Trial-by-Trial Basis

One of the characteristics of the data presented so far is that there is a very high correlation between the number of trials to criterion and the various measures of performance. Although the exact correlation coefficients are unnecessary in terms of the following argument, it may be noted that they range from .60 to .90 in most cases.

This high correlation is to be expected because the more trials one flies the more errors one is likely to make. Consequently measures such as "total altitude errors" or "total critical course errors" are refinements of the original "trials to criterion." What they do is bring a little greater sensitivity to the analysis under the assumption that even on trials where the criteria have been violated the better pilots will register fewer errors than those pilots who are experiencing problems.

The analysis in this section examines this within-trial performance across all the groups on those trials common to all subjects. This allows one to compare performances at defined points in the learning process rather than across whole sections of the learning process. The hypothesis is that differences should be apparent within trials as well as across trials.

Seven trials were common to all subjects in the holding pattern sequence. These were the first trial in the "test" sequence; and the first two trials in each of the three other sequences. What this means is that with respect to the "test" patterns, some subjects completed only one pattern before getting lost, while with respect to the other conditions at least one subject in each condition attained criterion in the minimum of two trials.

Figures 40 and 41 give the mean number of errors and critical errors by trial for all groups for each of the trials listed above. The solid lines join trials in the same condition, while the dotted lines join the different conditions for each group.

Analyses of variance were performed for errors and critical errors on each of the trials reported. Significant treatment differences occurred for critical errors in the test pattern trial (p = .035), and for errors in the first of the crosswind trials (p = .004). These are reported in Tables 27 and 28. A significant difference (p = .004) by order occurred for critical errors on the second trial of the crosswind condition (Table 29), although this was not in the direction hypothesized.

In this trial the groups that had extra practice in the form of the procedure turn trial performed worse than their counterparts.

In most of the analyses of variance, differences by treatment or group were very small, and any trends were weaker than those in the "trials to criterion" analyses.

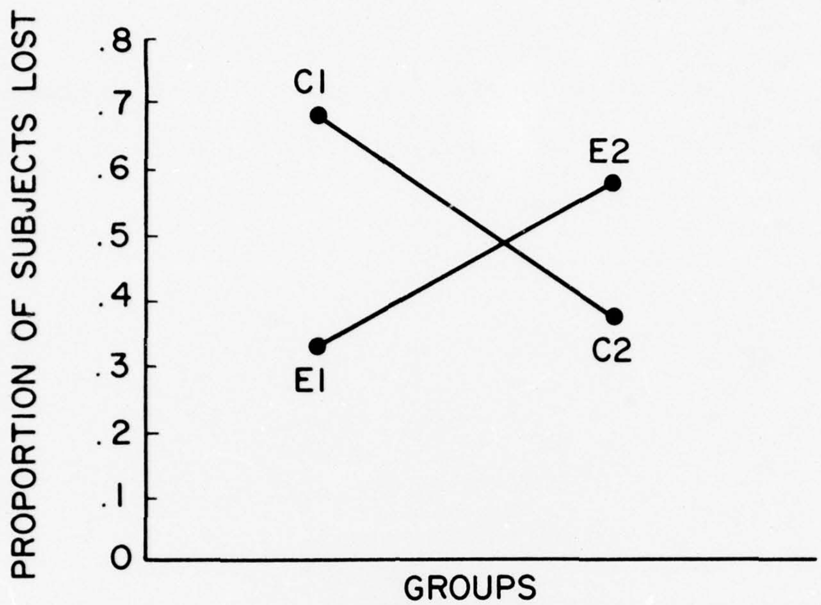

C1	.69
C2	.38
E1	.36
E2	.58

Figure 22. Proportion of subjects lost by group on "test" patterns.

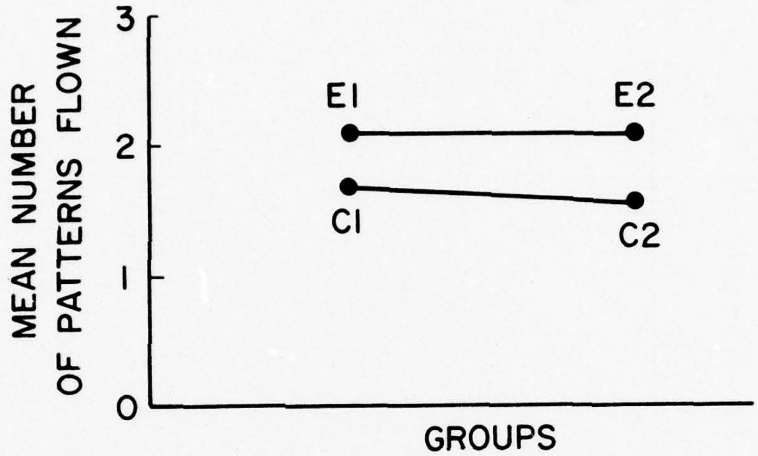

C1	1.69
C2	1.54
E1	2.18
E2	2.15

Figure 23. Mean number of patterns flown by group on "test" patterns.

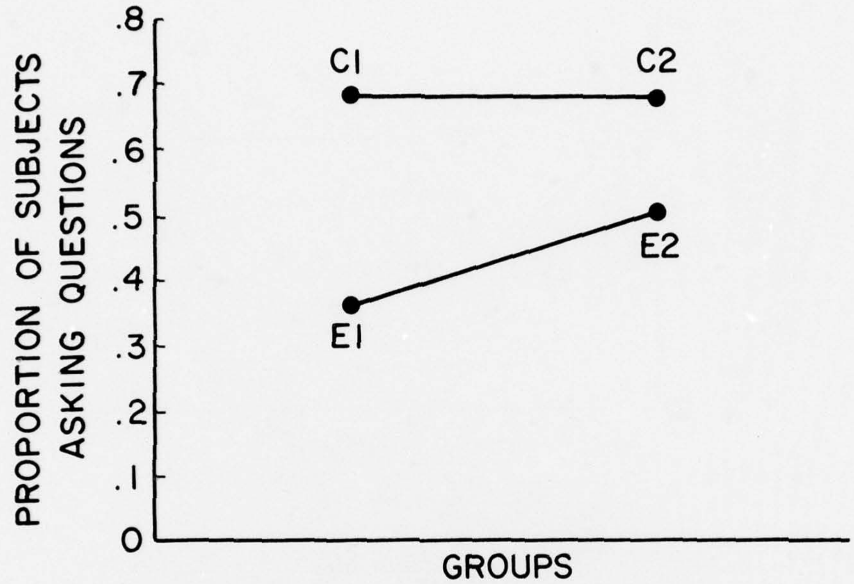

C1	.69
C2	.69
E1	.36
E2	.50

Figure 24. Proportion of subjects asking questions before "test" patterns.

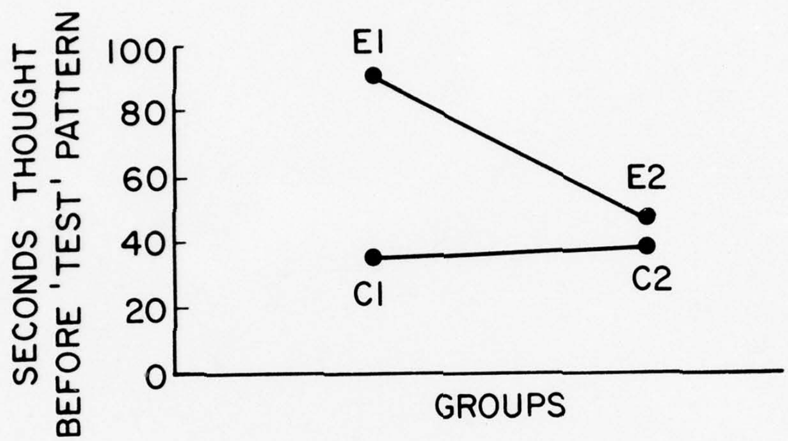

C1	37.5
C2	39.8
E1	94.2
E2	45.8

Figure 25. Time spent thinking before "test" patterns.

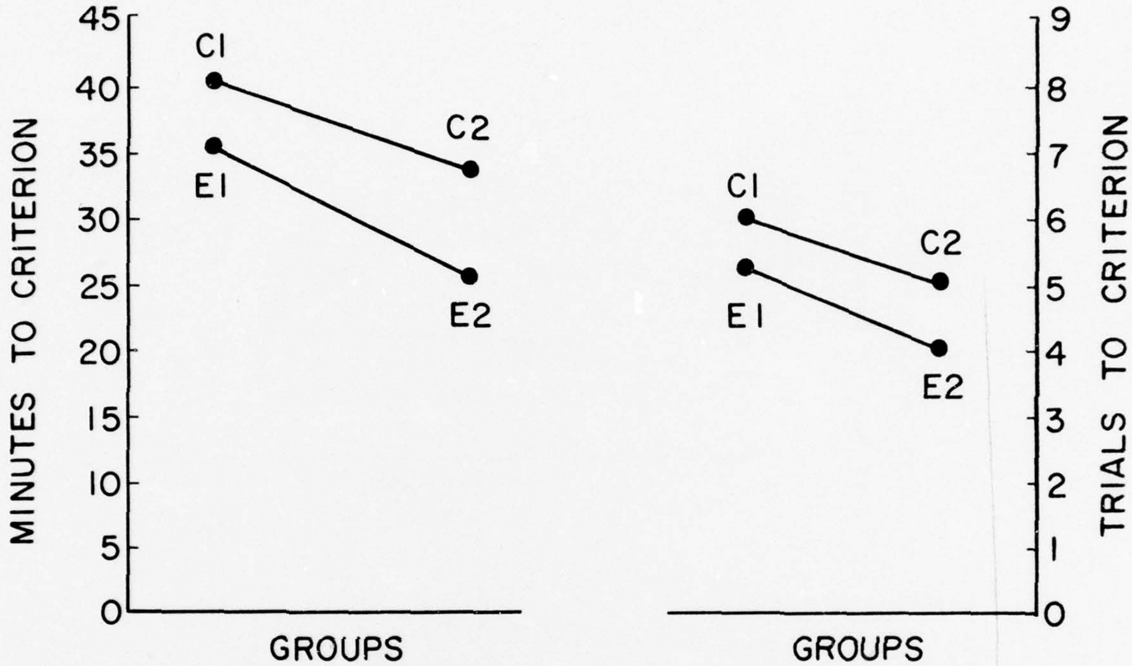

C1	40.5
C2	34.3
E1	36.4
E2	26.5

C1	6.0
C2	5.0
E1	5.4
E2	4.1

Figure 26. Time and trials to criterion on the no-wind pattern test.

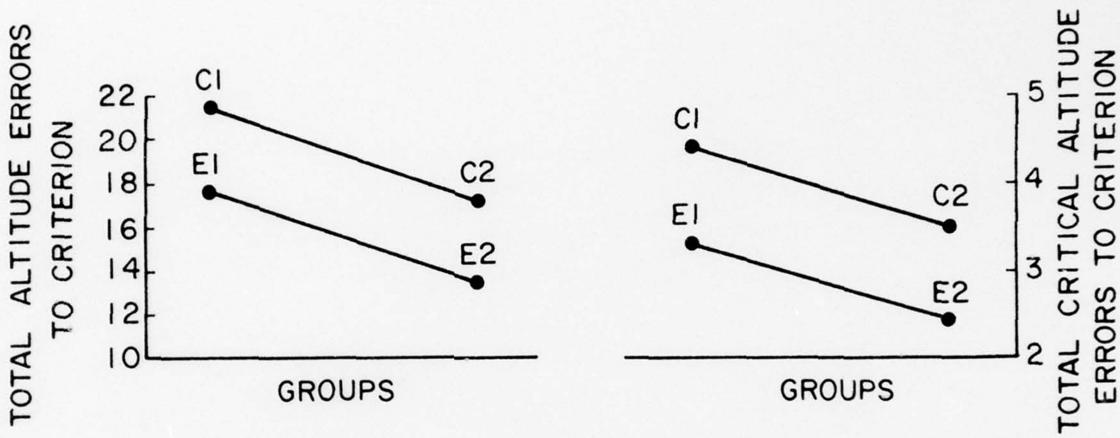

C1	21.3
C2	17.0
E1	17.5
E2	13.5

C1	4.3
C2	3.4
E1	3.2
E2	2.4

Figure 27. Altitude errors on the no-wind pattern task.

AD-A043 246 ILLINOIS UNIV AT URBANA-CHAMPAIGN SAVOY AVIATION RES--ETC F/G 5/9
AN EVALUATION OF A COMPLEX COMPUTER-BASED FLIGHT PROCEDURES TRA--ETC(U)
JAN 77 S R TROLLIP F44620-76-C-0009
UNCLASSIFIED ARL-77-1/AFOSR-77-1 AFOSR-TR-77-1013 NL

2 OF 2
AD
A043246

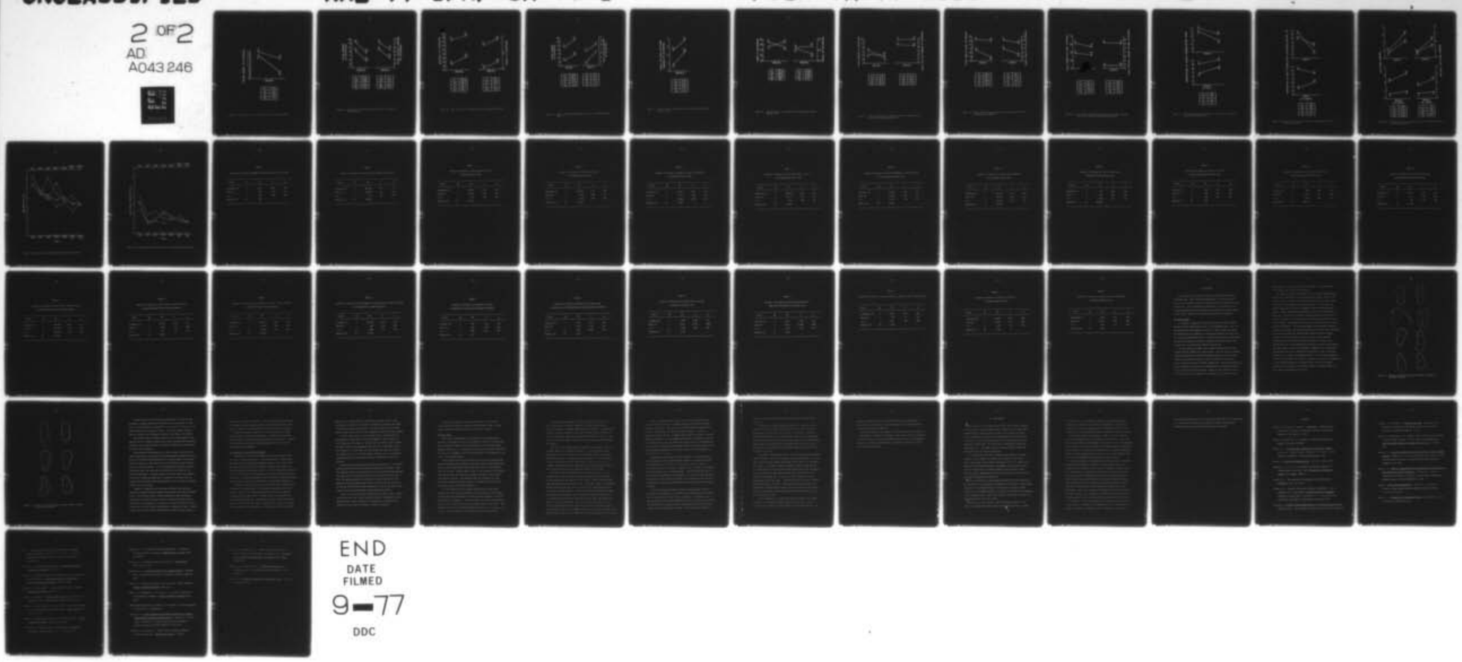

END
DATE
FILMED
9-77
DDC

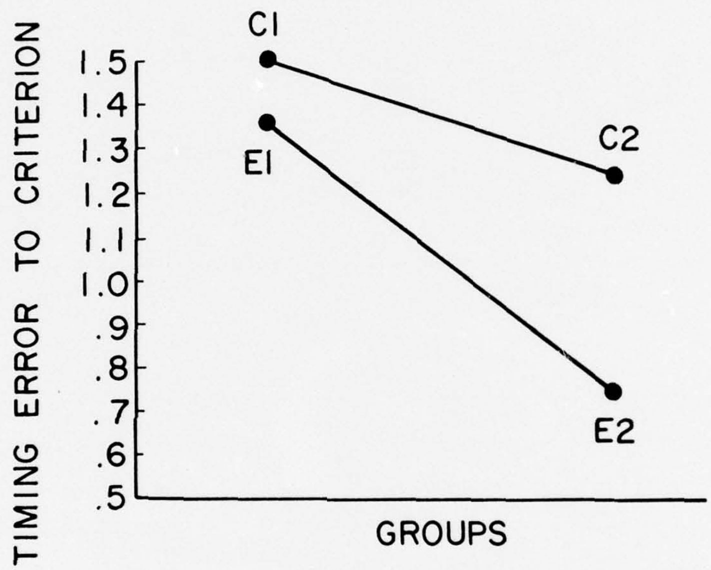

C1	1.50
C2	1.23
E1	1.36
E2	.73

Figure 28. Timing errors to criterion on the no-wind pattern task.

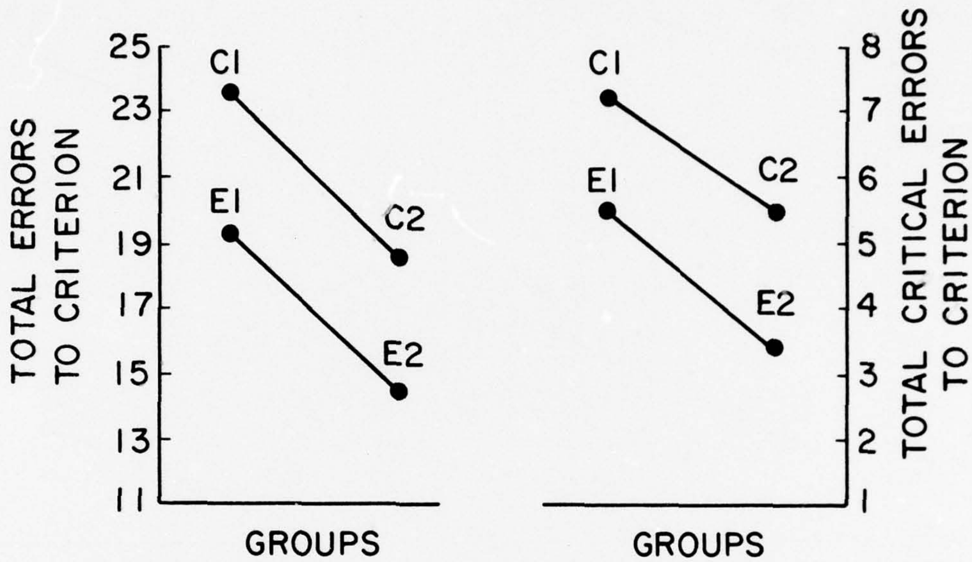

C1	23.3
C2	18.6
E1	19.3
E2	14.6

C1	7.1
C2	5.5
E1	5.5
E2	3.4

Figure 29. Total errors and total critical errors to criterion on no-wind task.

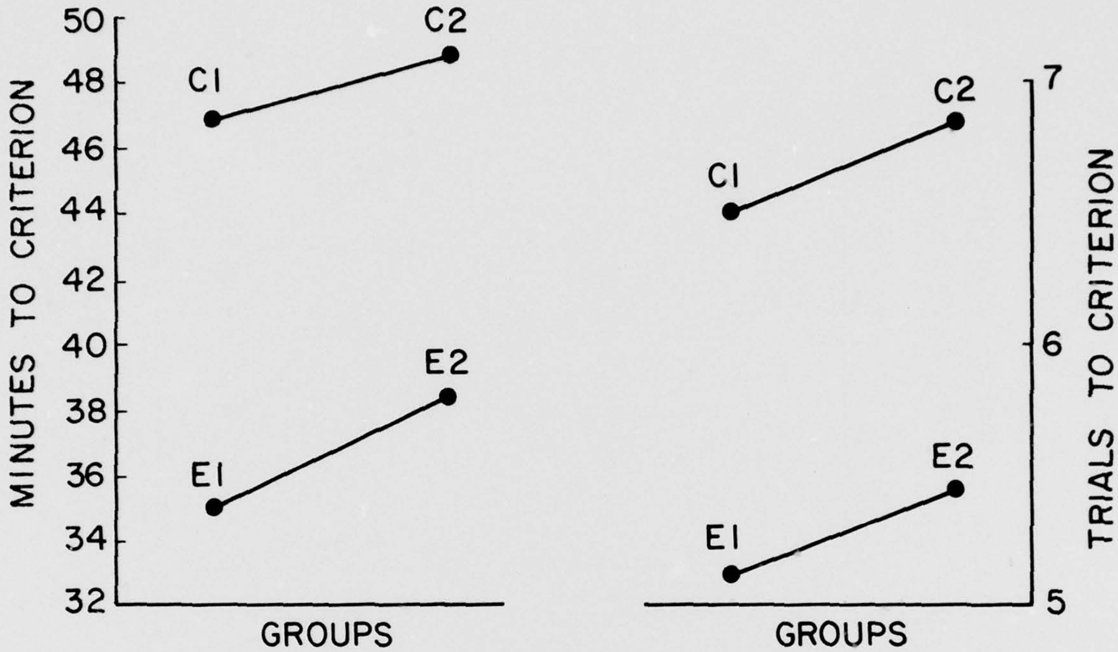

C1	46.9
C2	48.9
E1	35.1
E2	38.1

C1	6.5
C2	6.9
E1	5.1
E2	5.5

Figure 30. *Time and trials to criterion on the crosswind pattern task.*

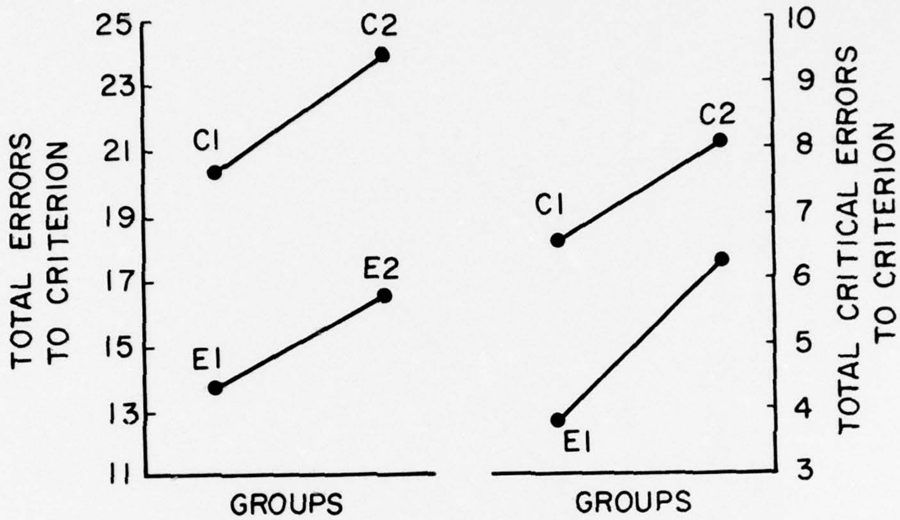

C1	20.5
C2	23.7
E1	13.8
E2	16.9

C1	6.6
C2	8.0
E1	3.7
E2	6.2

Figure 31. Total errors and total critical errors on crosswind pattern task.

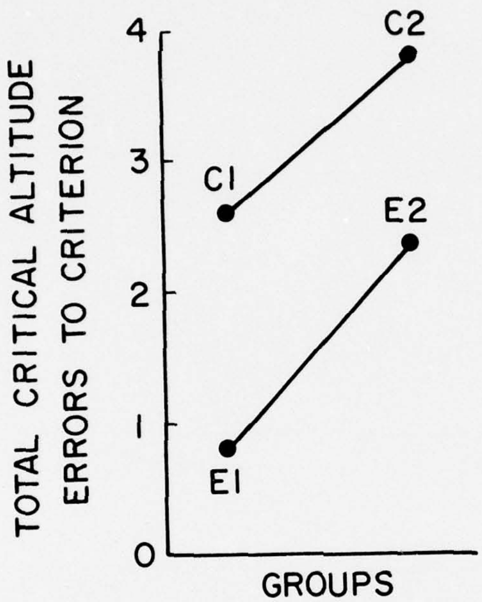

C1	2.5
C2	3.9
E1	0.8
E2	2.5

Figure 32. Total critical altitude errors to criterion on crosswind pattern task.

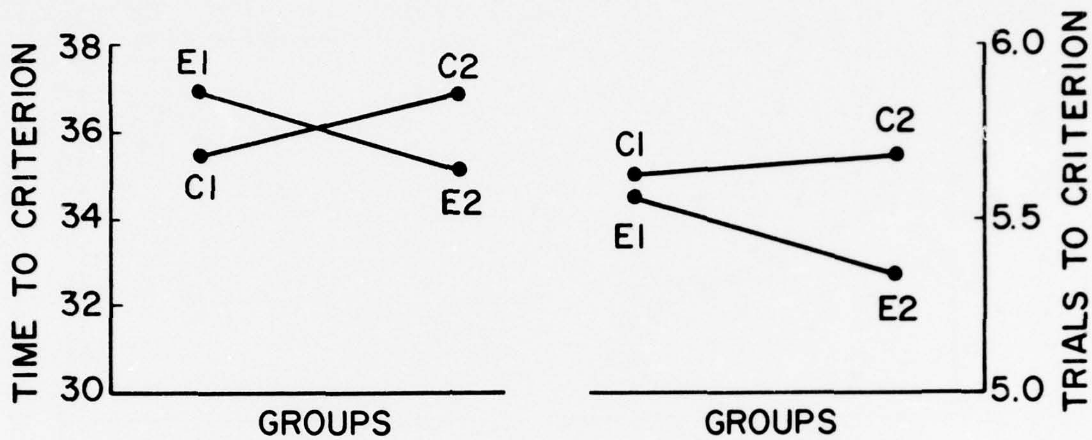

C1	35.5
C2	36.9
E1	36.9
E2	35.1

C1	5.62
C2	5.69
E1	5.55
E2	5.31

Figure 33. Time and trials to criterion on the quartering tailwind pattern task.

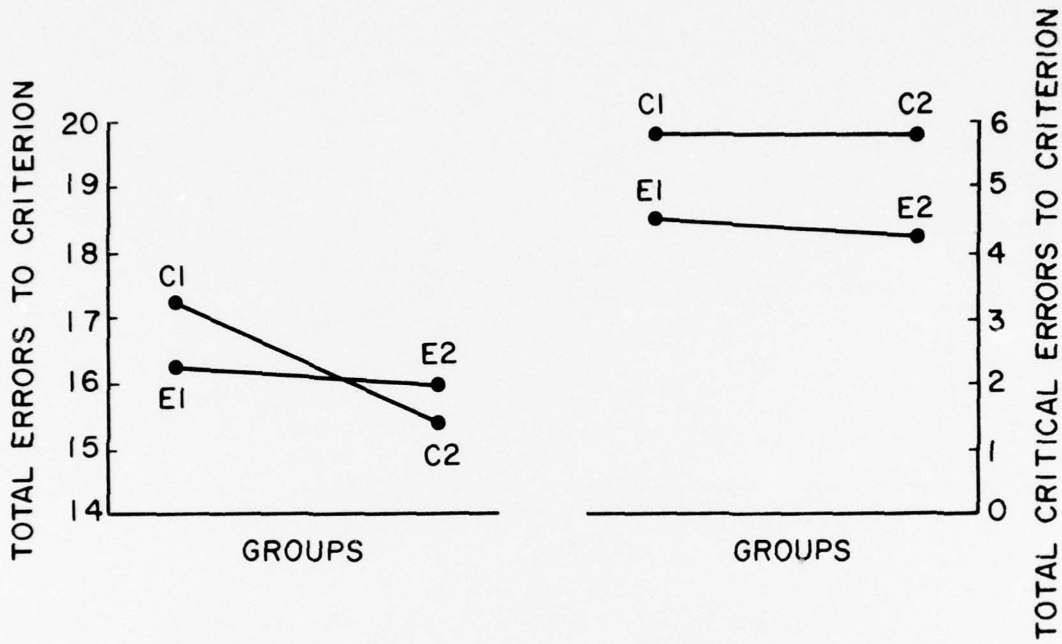

C1	17.2
C2	15.4
E1	16.2
E2	16.0

C1	5.7
C2	5.7
E1	4.5
E2	4.2

Figure 34. Total errors and total critical errors to criterion on the quartering tailwind pattern task.

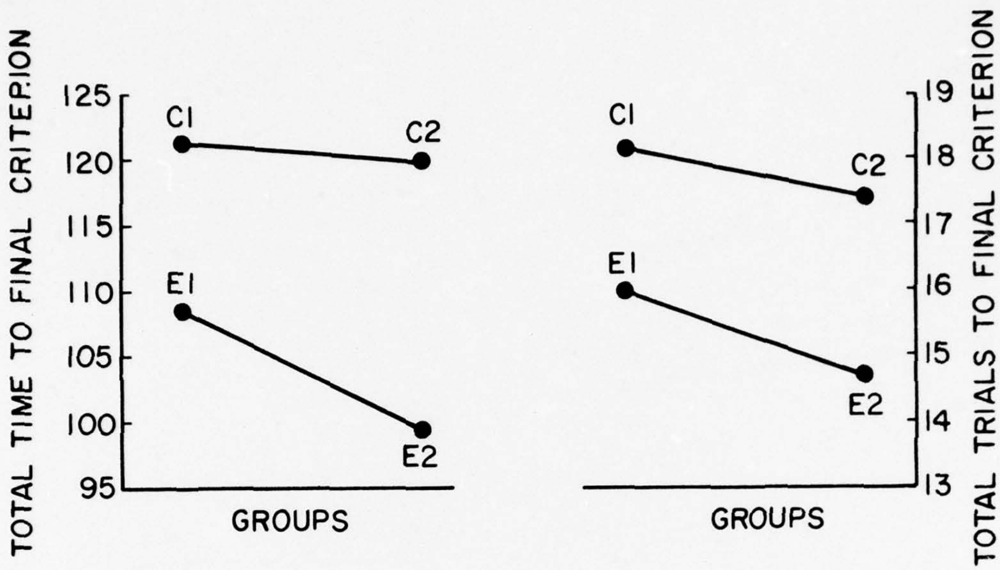

C1	122.9
C2	120.1
E1	108.4
E2	99.7

C1	18.1
C2	17.5
E1	16.0
E2	14.8

Figure 35. Total time and total trials to final criterion over entire holding pattern sequence.

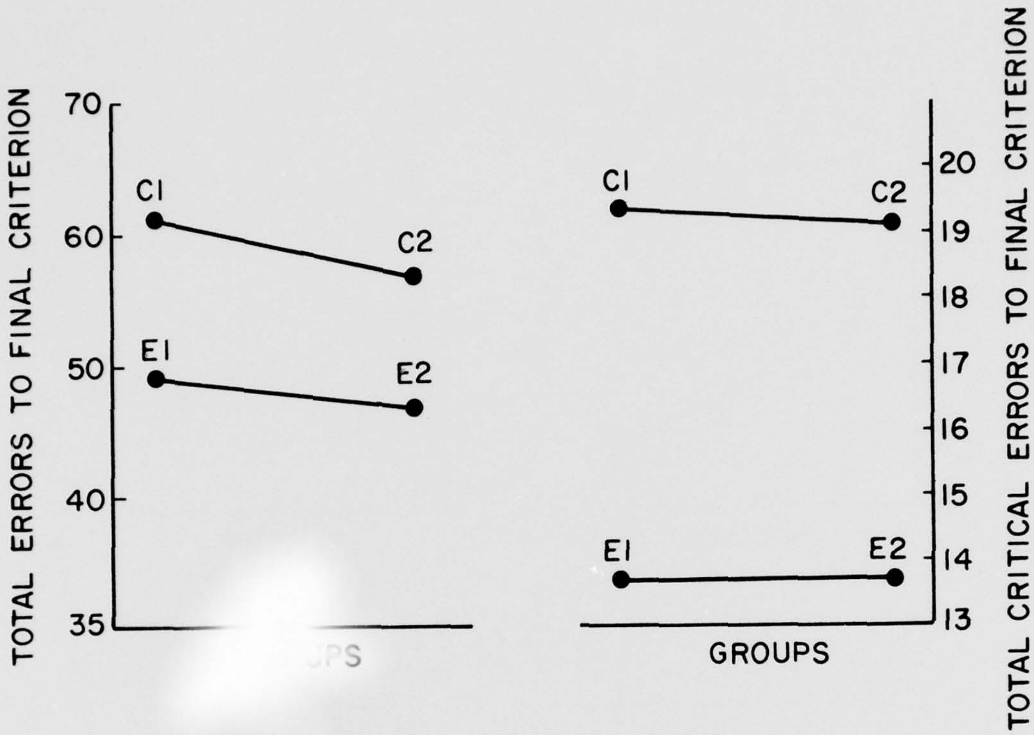

C1	61.1
C2	57.7
E1	49.3
E2	47.5

C1	19.5
C2	19.2
E1	13.7
E2	13.8

Figure 36. Total errors and total critical errors to final criterion over entire holding pattern training sequence.

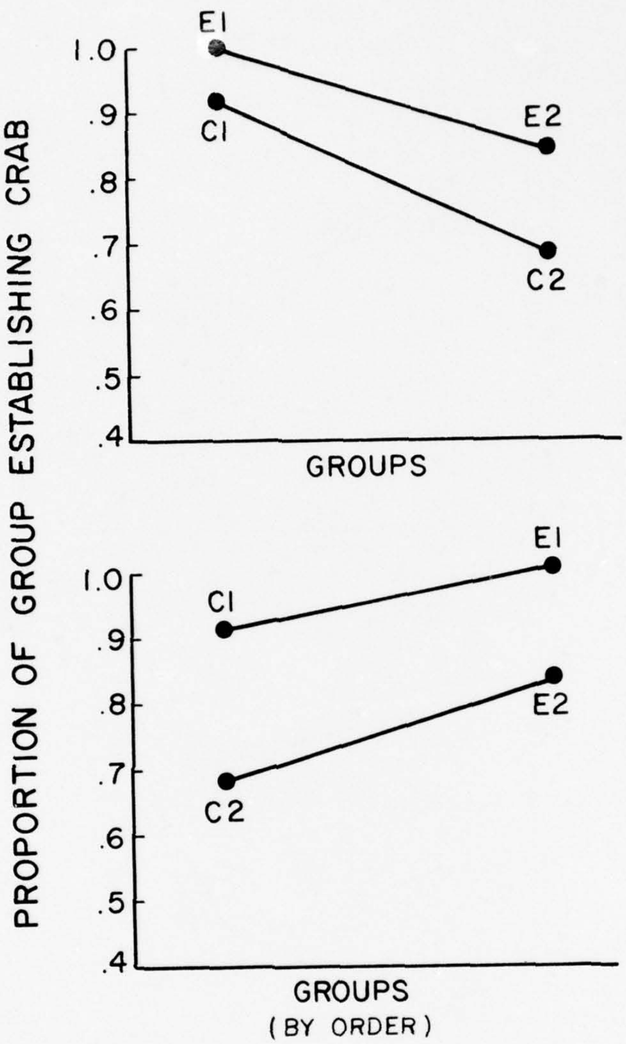

C1	.92
C2	.69
E1	1.00
E2	.85

Figure 37. Proportion of group establishing a crab before reaching OM outbound on procedure turn.

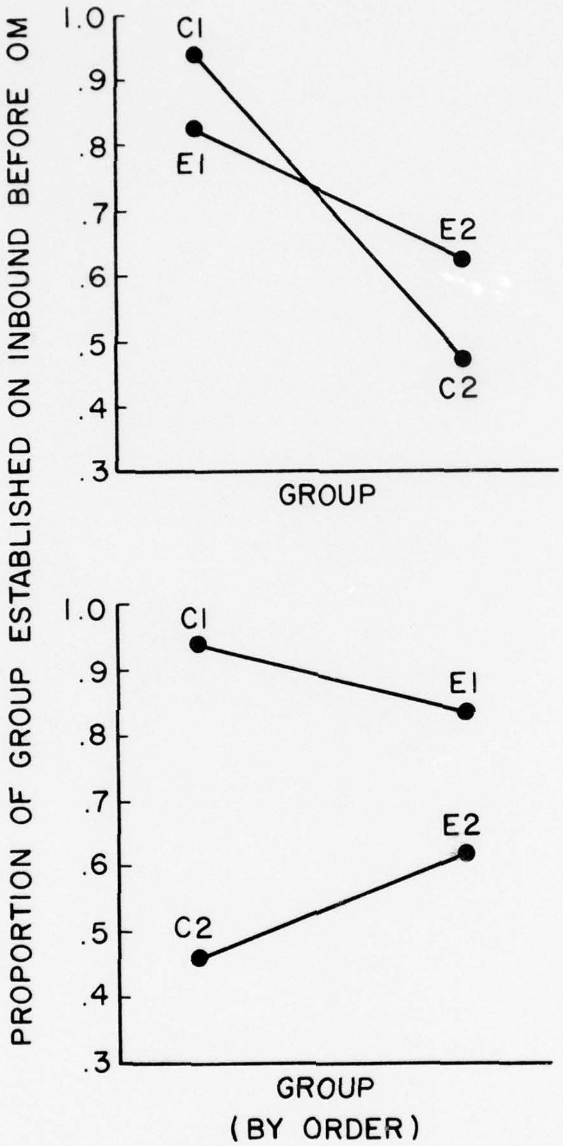

C1	.92
C2	.46
E1	.82
E2	.61

Figure 38. Proportion of groups established on inbound before OM by treatment and order.

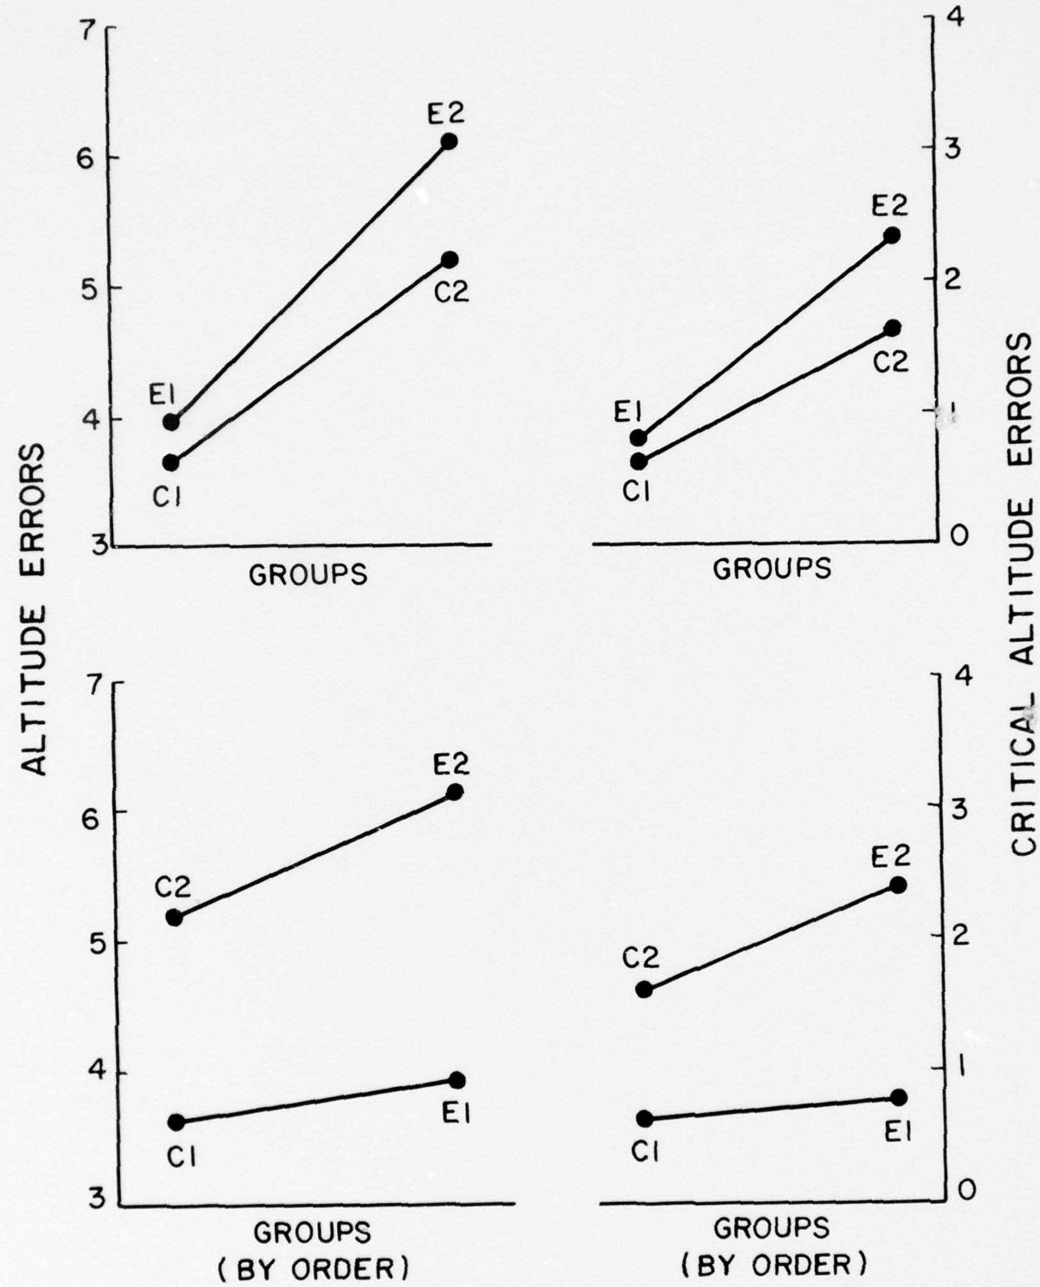

C1	3.62
C2	5.27
E1	3.91
E2	6.23

C1	.65
C2	1.62
E1	.73
E2	2.42

Figure 39. Altitude and critical altitude errors by treatment and order on procedure turn task.

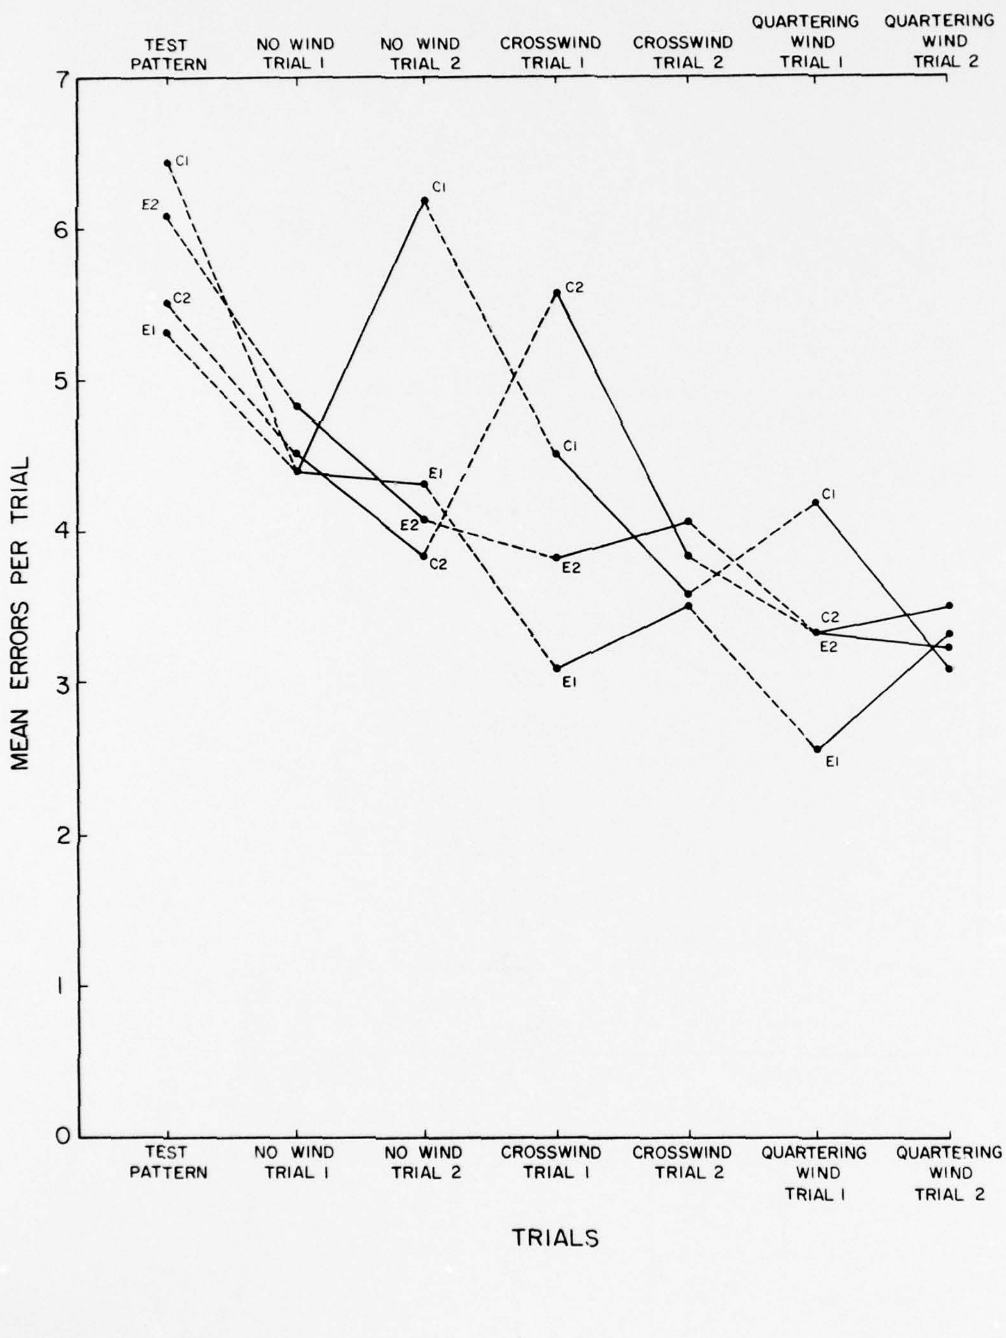

Figure 40. Mean errors per trial on the holding pattern sequence.

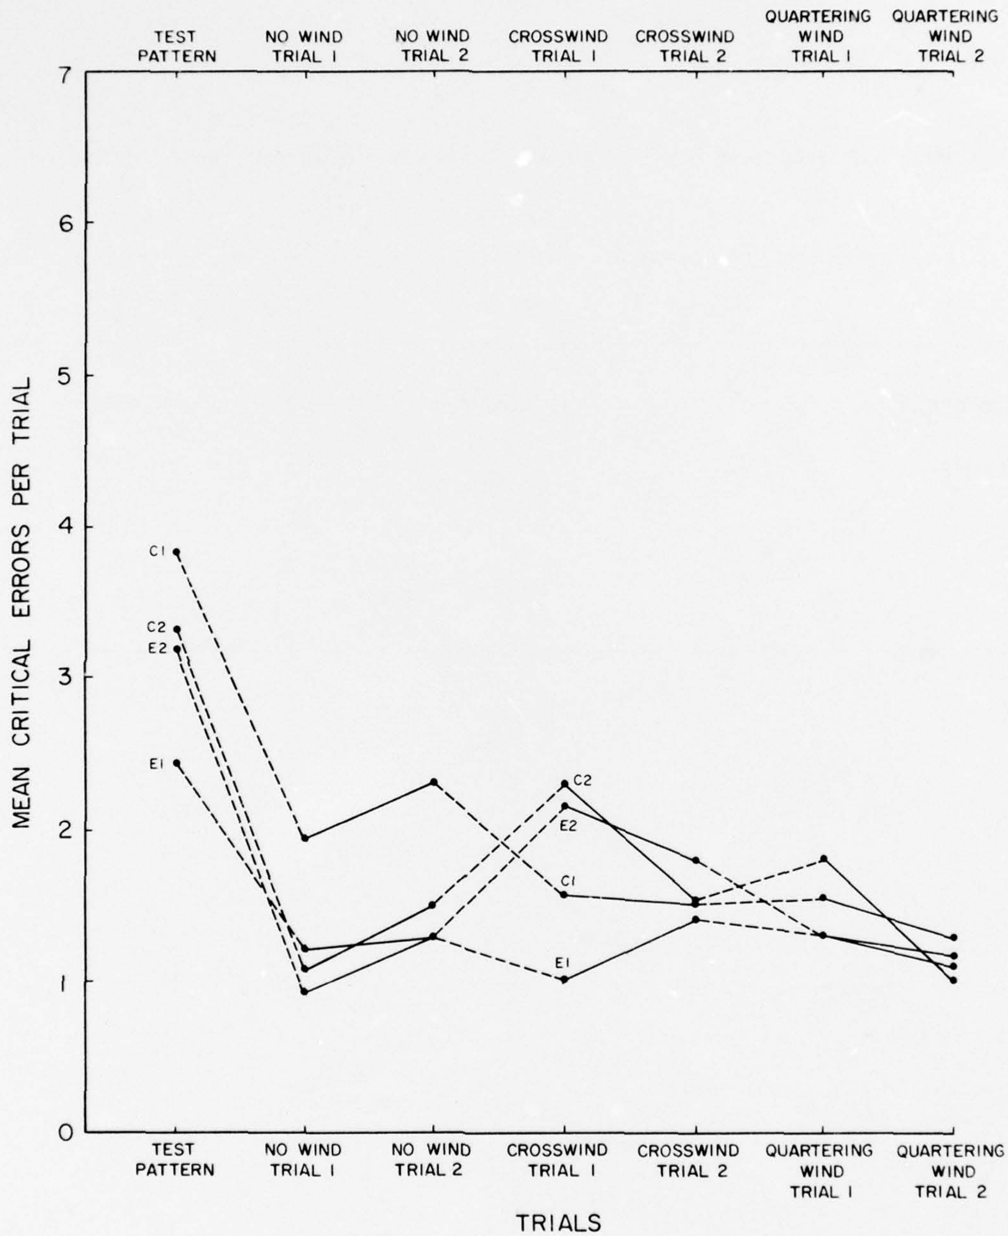

Figure 41. Mean critical errors per trial on the holding pattern sequence.

TABLE 7

Analysis of Variance for Number of Patterns Flown in "Test" Trial

Source	df	MS	F	p
Treatments (T)	1	3.795	3.851	.056
Order (O)	1	.102	.104	.748
T x O	1	.049	.050	.824
Subjects (S)	46	.985		

TABLE 8

Analysis of Variance of Total Errors on No-Wind Pattern Task

Source	*df*	MS	*F*	*p*
Treatments (T)	1	202.625	1.05	.311
Order (O)	1	278.498	1.44	.236
T x O	1	.017	.00	.992
Subjects (S)	46	192.940		

TABLE 9

Analysis of Variance of Total Critical Errors

on No-Wind Pattern Task

Source	df	MS	F	p
Treatments (T)	1	42.381	1.756	.192
Order (O)	1	40.634	1.684	.201
T x O	1	.662	.027	.864
Subjects (S)	46	24.131		

TABLE 10

Analysis of Variance for "Time to Criterion"

for Crosswind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	1577.203	4.211	.046
Order (O)	1	77.283	.206	.652
T x O	1	3.022	.008	.929
Subjects (S)	46	374.548		

TABLE 11

Analysis of Variance for Number of Trials to Criterion

for Crosswind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	23.599	2.693	.108
Order (O)	1	1.773	.202	.655
T x O	1	0.0006	.0001	.993
Subjects (S)	46	8.762		

TABLE 12

Analysis of Variance for the Total Number of Errors

on the Crosswind Pattern Task

Source	*df*	MS	*F*	*p*
Treatment (T)	1	568.910	4.068	.049
Order (O)	1	123.275	.881	.353
T x O	1	.024	.0002	.989
Subjects (S)	46	139.861		

TABLE 13

Analysis of Variance for the Total Number of Critical Errors on the Crosswind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	68.753	3.335	.074
Order (O)	1	46.238	2.243	.141
T x O	1	4.214	.204	.653
Subjects (S)	46	20.613		

TABLE 14

Analysis of Variance for "Total Altitude Errors"

for Crosswind Pattern Task

Source	*df*	MS	*F*	*p*
Treatment (T)	1	441.672	3.634	.063
Order (O)	1	123.549	1.017	.319
T x O	1	14.468	0.119	.732
Subjects (S)	46	121.525		

TABLE 15

Analysis of Variance for Time to Criterion on

Quartering Tailwind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	.703	.003	.961
Order (O)	1	.623	.002	.963
T x O	1	32.168	.113	.738
Subjects (S)	46	283.826		

TABLE 16

Analysis of Variance for Trials to Criterion

on Quartering Tailwind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	.642	.113	.738
Order (O)	1	.080	.014	.906
T x O	1	.308	.054	.817
Subjects (S)	46	5.681		

TABLE 17

Analysis of Variance for Total Errors to Criterion on Quartering Tailwind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	.566	.006	.937
Order (O)	1	12.873	.146	.705
T x O	1	7.027	.079	.779
Subjects (S)	46	88.491		

TABLE 18

Analysis of Variance for Total Critical Errors on Quartering Tailwind Pattern Task

Source	df	MS	F	p
Treatment (T)	1	23.719	1.652	.205
Order (O)	1	.477	.033	.856
T x O	1	.477	.033	.856
Subjects (S)	46	14.359		

TABLE 19

Analysis of Variance for Total Time to Final Criterion on Entire Holding Pattern Training Sequence

Source	df	MS	F	p
Treatment (T)	1	3795.988	2.802	.101
Order (O)	1	412.376	.304	.584
T x O	1	105.486	.078	.781
Subjects (S)	46	1354.851		

TABLE 20

Analysis of Variance for Total Trials to Final Criterion
on Entire Holding Pattern Training Sequence

Source	df	MS	F	p
Treatment (T)	1	70.709	2.632	.112
Order (O)	1	8.903	.331	.568
T x O	1	1.177	.044	.835
Subjects (S)	46	26.866		

TABLE 21

Analysis of Variance for Total Number of Errors to Final Criterion on Holding Pattern Training Sequence

Source	*df*	MS	*F*	*p*
Treatment (T)	1	1508.437	3.030	.088
Order (O)	1	84.150	.169	.683
T x O	1	6.897	.014	.907
Subjects (S)	46	497.904		

TABLE 22

Analysis of Variance for Total Number of Critical Errors to Final Criterion on Holding Pattern Training Sequence

Source	*df*	MS	*F*	*p*
Treatment (T)	1	239.438	6.736	.013
Order (O)	1	.682	.019	.890
T x O	1	2.689	.076	.785
Subjects (S)	46	35.546		

TABLE 23

Analysis of Variance on Proportions of Groups

Establishing a Crab before Reaching OM Outbound

Source	df	MS	F	p
Treatment (T)	1	.166	1.414	.240
Order (O)	1	.460	3.929	.053
T x O	1	.018	0.157	.694
Subjects (S)	46	.117		

TABLE 24

Analysis of Variance on Proportions of Groups being Established on the Inbound Radial before Reaching OM Inbound

Source	*df*	MS	*F*	*p*
Treatment (T)	1	.007	.039	.845
Order (O)	1	1.372	7.118	.011
T x O	1	.208	1.080	.304
Subjects (S)	46	.193		

TABLE 25

Analysis of Variance for Altitude Errors after OM Outbound on Procedure Turn

Source	*df*	MS	*F*	*p*
Treatment (T)	1	4.898	1.569	.217
Order (O)	1	49.132	15.739	.0003
T x O	1	1.386	0.444	.508
Subjects (S)	46	3.122		

TABLE 26

Analysis of Variance for Critical Altitude Errors

Made after OM Outbound on Procedure Turn

Source	df	MS	F	p
Treatment (T)	1	2.412	1.738	.194
Order (O)	1	21.952	15.809	.0003
T x O	1	1.676	1.207	.278
Subjects (S)	46	1.389		

TABLE 27

Analysis of Variance for Critical Errors on the First "Test" Pattern Trial

Source	*df*	MS	*F*	*p*
Treatments (T)	1	8.034	4.72	.035
Order (O)	1	.213	.13	.725
T x O	1	5.277	3.10	.085
Subjects (S)	46	1.704		

TABLE 28

Analysis of Variance for Errors on the First

Crosswind Condition Trial

Source	df	MS	F	p
Treatments (T)	1	32.168	8.94	.004
Order (O)	1	10.435	2.90	.095
T x O	1	.703	.20	.661
Subjects (S)	46	3.598		

TABLE 29

Analysis of Variance for Critical Errors on the First Crosswind Condition Trial

Source	*df*	MS	*F*	*p*
Treatments (T)	1	1.840	1.64	.207
Order (O)	1	10.595	9.46	.004
T x O	1	.662	.59	.446
Subjects (S)	46	1.120		

9. DISCUSSION

The discussion of the results follows the same structure used in presenting them. First, the PLATO program and its acceptance by the subjects are considered; second, the reliability of the scoring procedure is discussed. This is followed by a discussion of the effects of the training treatments on the performance of the procedure turn and holding pattern tasks. Finally the trial-by-trial results are considered.

The PLATO Program

A fundamental assumption of this entire study was that the PLATO package actually taught students how to fly holding patterns. This was incorporated into the program in the form of PLATO's flight instructor who determined on the completion of each pattern whether the pattern had met certain predetermined criteria. If these criteria were satisfactorily attained the student was allowed to progress to the next stage in the program, and on the same basis, finally reach the end.

Had the package not taught, these criteria would never have been attained (except possibly on a chance basis). The fact that all students finally "graduated" from the PLATO training package is in itself proof either that learning took place or that the tasks were so easy that everyone could pass without learning anything new. This second possibility may be discarded on the basis of an examination of the patterns flown by various subjects during the program. Figures 42 and 43 show the series of patterns flown by two subjects in attaining final criterion on the

PLATO program. In both cases the series exhibits a strong learning effect within each of the wind conditions.

In Figure 42, Frame 1, the subject flew this no-wind pattern quite well until on inbound when he over-corrected for being a little to the right of course. In Frame 2, he recovered well from a bad start, where his initial approach must have been very erratic. Frame 3 shows the subject making one of the most fundamental errors in holding pattern flying. Having established the appropriate crab on the inbound leg (in this case to the left) he crabbed in the wrong direction (once again to the left, instead of to the right) on the outbound and was blown a long way off course. His attempt to fly back to the inbound radial was futile in the circumstances. On his second attempt in the crosswind condition, (Frame 4) he corrected his previous mistake, but failed to crab sufficiently into the wind to make much progress back towards the inbound radial. Finally, on his third attempt, he flew an almost perfect pattern. Progressing to the quartering headwind condition, the subject took insufficient crab on the outbound and was blown across the inbound radial, and also failed to shorten the outbound to compensate for the tailwind, and was blown far from the holding fix, resulting in about a 90-second inbound leg. In Frame 7 he compensated well for the crosswind component, but still flew outbound for too long. Finally, in Frame 8, his compensation for the headwind component was adequate and he flew a good pattern. Learning is also evident in the second sequence of drawings (Figure 43), where similar improvements can be seen.

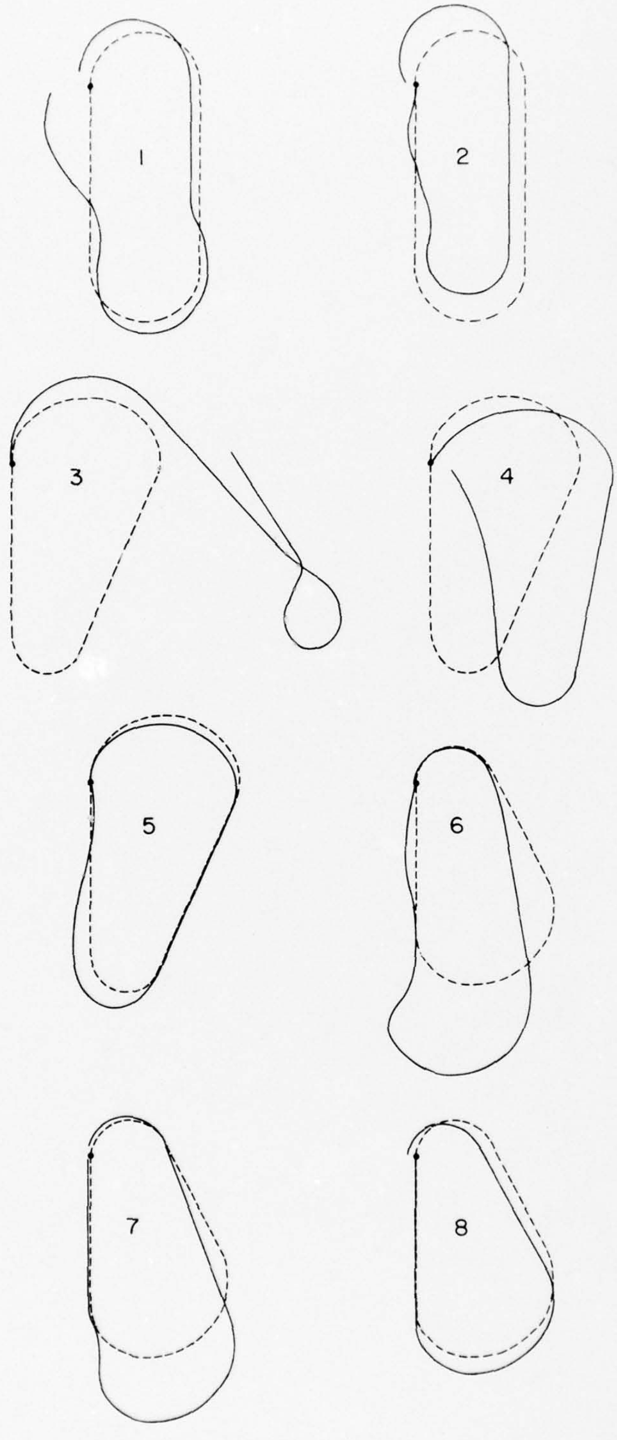

Figure 42. Sequence of patterns flown by one subject on PLATO in attaining criterion.

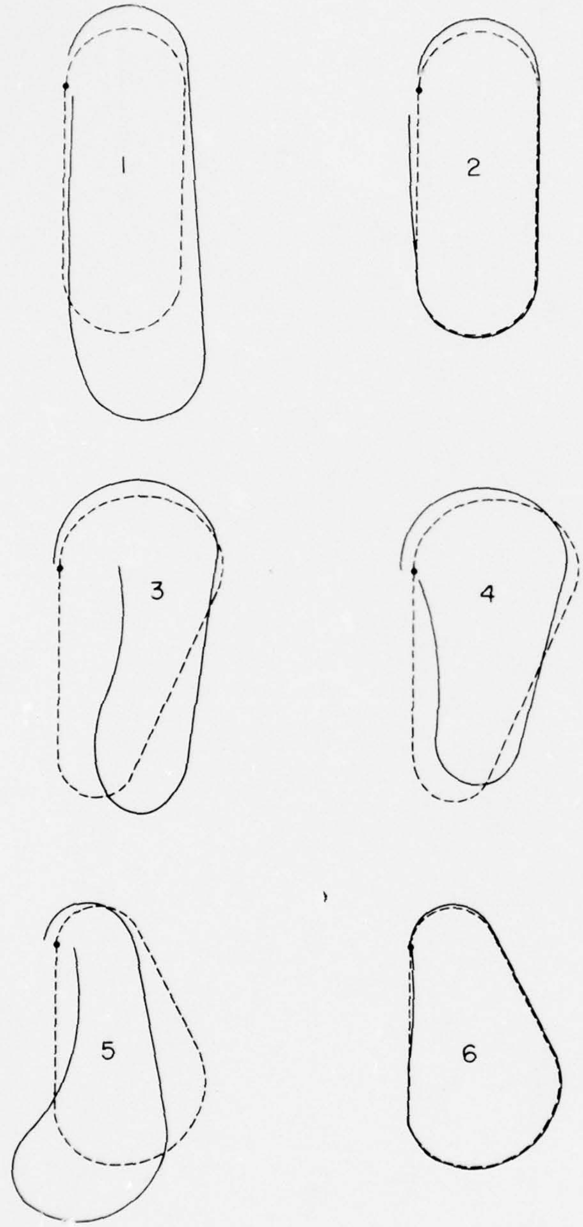

Figure 43. Sequence of patterns flown by another subject on PLATO in attaining criterion.

Together with the high correlations between time on PLATO and some performance variables measured in the GAT-2, the learning shown in these (and all other) sequences of patterns gives considerable credibility to the PLATO package and its capacity to teach. In their own right, these are important results, and by themselves justify the writing of the program.

This justification is further enhanced by the attitude shown by the subjects who were trained on PLATO, many of whom, spontaneously and without prompting, gave as their reasons for liking the program those very reasons that led to its initiation.

Some subjective observations are in order; several subjects did not spend as much time on the section of the program which allowed them to see what effect different wind conditions had on pattern shape as the experimenter had expected. Some tried only one wind condition before opting to proceed to the next section. It is the experimenter's opinion that those subjects who spent the least time on this section also ended up being the worst performers. A possible reason for this was that these subjects had not yet realized the importance of working out in advance what the pattern would look like, and so when they had to fly a pattern they started without adequate planning.

Secondly, it appeared that the better performers also learned to manipulate the hand-controller quicker than those who eventually performed worse. This would have been a possible problem had these differences disappeared in the GAT-2 trials, because it could have implied that bad performance was contingent on an inability to operate the hand-controller, rather than on ability levels closer related to instrument flying. However, because those subjects who had difficulty with PLATO's patterns tended to

experience difficulty in the GAT-2, it was concluded that the operation of the hand-controller depended somewhat on flying ability rather than ability to operate the hand-controller influencing flight performance.

Finally, the program worked well on a stand alone basis, without needing someone to supervise each student. Ultimately one could have PLATO terminals and hand-controllers available for use any time. Overall the PLATO program achieved its goals pedagogically, and did so with considerable user enthusiasm.

The Reliability of the Scoring Procedure

Fundamental to the scoring procedures were scoring booklets that were easy to understand and easy to mark appropriately. Other studies that used similar equipment and procedures (Koonce, 1974; Jacobs, 1976) have shown that the use of two independent observers using relatively objective measuring instruments leads to high reliability between observations. That this is true has obviated the need for expensive electronic recording and deciphering equipment, whose reliablity is sometimes questionable, and whose results are often difficult to interpret meaningfully.

The observers in this study sat one behind the other minimizing the opportunity of collaboration. In some cases disagreements arose as to whether a subject had attained the necessary criteria, one observer scoring "yea," the other "nay." In these cases it was presumed that criterion had not been met. Discrepancies like this could arise from a number of factors, the most likely of which was the effects of parallax on the reading of the instruments. The instructor sitting in the right

front seat of the GAT-2 observes the flight instruments from an acute angle making it difficult to determine in marginal cases where the needle or pointer in question actually was. The second observer, on the other hand, sitting further back, saw the instruments from a more normal position.

A second common source of error between the two observers was the determination of just when a rollout had stopped, or when straight and level flight had been achieved. In some cases the subject would diminish his angle of bank at an appropriate time, but fail to establish a wing level attitude. In such cases the two observers may start the timing of the inbound leg, say, at different times, resulting in one recording that criterion had been attained while the other recorded that it had been violated.

There also was some correlation between the importance of the measure being recorded and the observer-observer reliability coefficient. What this suggests is that the observers were less attentive when they realized that their lack of attention would not affect the outcome of the trial. Thus, the observer-observer reliability for measures recorded on the portion of the flight between the initial starting point six miles from the VOR and the VOR itself (which signalled the start of the pattern) was lower than for the same measures recorded during the pattern.

Apart from these factors that affected the scoring, careful training and practice took care of other potential problems. Ambiguities in the procedures were either eliminated or explained, and consequently the observers did not report any trouble in using the booklets.

It must be concluded, as Koonce and Jacobs did, that the use of a well-constructed relatively objective scoring booklet tends to result in high correlations between the scores of trained observers.

The GAT-2 Tasks

As has been mentioned before, the purpose of the procedure turn trial was to determine if there was any positive transfer from the holding pattern instruction both prior to the completion of the holding pattern sequence, and after it. It turned out that if any such transfer took place its effects were swamped by the very strong effects of the position of the turn in the training sequence.

Those subjects who flew the procedure turn after the holding pattern sequence performed almost perfectly on the main task of establishing themselves on the inbound radial, while those who flew it at the beginning fared not too well. However, when it came to the holding pattern sequence those who had flown the turn first generally performed a little better than those who had not. This indicates that the procedure turn when flown first provided additional practice in flying the GAT-2 which was detectable later in the training sequence.

It also indicates that the initial familiarization ride, while demanding the same level of performance from all subjects, may not have provided sufficient experience to prevent increased exposure to the GAT-2 from improving performance. Thus on attaining criterion on the familiarization task the subjects had not yet mastered control of the GAT-2. More stringent criteria on the familiarization flight might have minimized this effect. Thus the procedure turn task did not satisfy its purpose.

The holding pattern sequence, on the other hand, worked out very well, satisfying the necessary conditions to justify being called a "training" sequence. Subjects learned from the experience, as demonstrated by their ultimate attainment of the final criteria, and the condition-to-condition performance improved as may be seen in Figures 40 and 41. Both these Figures show definite learning curves as subjects progressed through the sequence.

The "test" pattern which began the sequence was designed to give an immediate measure of treatment differences. It was hypothesized that giving such a difficult pattern at the outset would improve the chances that differences due to treatment would be observed. It was feared that if the no-wind pattern were the first task, it would be so easy that treatment differences would not show, and that it would provide a practice effect which would help dilute treatment effects on later tasks.

The "test" pattern achieved its goal, because there was an almost significant ($p = .056$) difference between the number of patterns flown by the groups before the instructor took control. Those subjects who had taken the PLATO program performed better. Two further items of information emerged from the "test" pattern data, which may be interpreted in a way to support the PLATO approach. After the instructor had played the taped instructions to the subject, which included the information that he was to fly patterns without knowing anything about the wind, each subject was asked if he had any questions. There was a strong indication ($p = .07$) that the experimental subjects asked fewer questions than those who had the traditional lecture. Furthermore, when told to take time to think about what they were going to do, the PLATO subjects thought longer about their task.

It may be argued that this was the case because the PLATO subjects had flown holding patterns on the PLATO program and thus had the unfair advantage of practice. This is true: the PLATO subjects definitely had practice, but this does not detract from the findings. In fact, it supports the whole idea that a computer-based system can provide a valuable learning experience. If, by spending time on PLATO, the subjects were able to understand better what was involved in flying patterns, and subsequently took more time to resolve these problems before starting the patterns, then the program had achieved one of its goals, namely to teach foresight, or thinking ahead of the plane.

Following the execution of the "test" pattern, each student entered the training part of the holding pattern sequence. Starting with the no-wind condition they progressed to the quartering tailwind condition, achieving criterion at each stage. Figures 40 and 41 support the assumption that this sequence was a training sequence. For both errors and critical errors definite improvements in performance occur as the subject progresses. Although every trial flown is not shown, one can easily see that the first two patterns of each wind condition were flown better than in the previous condition.

Overall performance through the entire training package indicates that those groups spending time on PLATO before entering the GAT-2 learned faster, and made fewer critical errors in doing so. Although on most of the performance measures these differences were not statistically different, the experimenter is satisfied that the data show sufficiently consistent trends for this conclusion to be drawn. The fact that the experimental

groups made significantly fewer critical errors overall supports this conclusion.

In any study of complex human performance it is always difficult to conclude with any certainty that what one thinks is causing improved performance is actually the correct cause. For instance, although the experimenter would like to conclude that the reason for the experimental groups tending to perform better was because exposure to the PLATO program caused them to build better mental representations of the task, and to have pictures of each pattern in their heads, other reasons may actually be closer to the mark.

For example, an observation made by all the instructors and observers deserves mention. Many of the subjects, most of whom had made use of the VOR in their previous training and flying, did not seem to be comfortable with its use. On several occasions a subject would be flying outbound in a right pattern (at which point the VOR needle should have been pegged to the left) when he would suddenly think that he had to turn left to intercept the inbound leg. Of course, in this situation the subject should have flown to the right. This difficulty was noticed in subjects both on PLATO and in the GAT-2. Consequently it is possible that the experimental groups performed better, not because they had better mental pictures, but because PLATO had afforded them the opportunity to learn how to interpret the VOR.

Such a conclusion, if it were true, would not detract from PLATO's ability to teach complex procedural tasks, but would be a cause to re-examine the fundamental skills required to fly holding patterns. If

complete confidence in the use of VOR's was all that was required to perform holding patterns well, then the emphasis in the PLATO program on mental pictures would be misplaced.

What this suggests is that while it is likely that the PLATO approach in emphasizing the dynamics of the shape of the pattern, and the opportunity it gave for subjects to practice patterns, were the causes for experimental differences, further experimentation is necessary with stricter control kept on entering abilities.

10. CONCLUSIONS

The results of this evaluation were better than expected. Whereas it had been hoped that the trends would be strong that the PLATO groups performed better, learned quicker and made fewer mistakes than their control counterparts, it had not been expected due to the difficulties of control of such a complex study that any of these trends would show statistical significance. The fact that the experimental groups made significantly ($P < .05$) fewer critical errors than their controls was added support for the original hypothesis.

It is surprising that such a relatively small period of training compared to total flight time (about three percent of total flight hours) could result in such marked differences. This is a measure of the strength of the PLATO instruction, and suggests that there may be many areas of training where an abstract approach which purposely avoids faithful physical simulation in favor of the simulation of the necessary cues, may be padagogically beneficial.

Part of the success of the computer approach used in this study may be attributed to the opportunity it provides for the student to participate actively in the learning process. By being able to do what is being learned in an environment conducive to learning, the student may concentrate on the important issues, rather than having to concentrate on factors not directly related to the task at hand.

Another part of the success of the PLATO approach may be in its ability to generate appropriate mental images. The fact that a student

"see" exactly how his performance compares with the ideal provides a capability not available in most other training modes. The wealth of information contained in these pictures is immediately apparent to the student, who is then encouraged to generate similar pictures in his mind.

Three further issues are now open to study, given the conclusion that the PLATO approach is at least tentatively successful. First, the exact cost and transfer effectiveness of the approach need to be evaluated; second, whether the success achieved by the PLATO program in transferring to the GAT-2 can be replicated in transferring to an airplane; and third, whether the PLATO approach may be used, not only to train neophytes, but also to retrain experienced pilots who are no longer current.

This latter point is of considerable interest, because if the PLATO experience is sufficiently like that of flying holding patterns in reality, then the approach could be used as a possible recurrency check. It is interesting to speculate into the future to the time when pilots could have relatively inexpensive micro-processors linked into their home television sets allowing them to fly various instrument maneuvers such as holding patterns or instrument approaches in comfort of their homes and at the same time getting appropriate feedback as to their performance. Quite often rustiness in instrument skills may be overcome by a brief exposure to an instrument environment, and such a computer-oriented approach as used in this study, may provide sufficient practice.

The results of this study indicate that complex procedural skills can be taught by a computer if the appropriate cues for performance are adequately simulated. Furthermore, such a simulation, even though it

may operate continuously as far as the user is concerned, can be implemented on a large-scale CAI system such as PLATO IV without there being any detrimental effects on the interactive nature of the system.

REFERENCES

Ableson, H., Goodman, N., Rudolf, L. Logo Manual. Cambridge, Mass.: Massachusetts Institute of Technology, Artificial Intelligence Laboratory, LOGO memo No. 7, 1974.

Alpert, D. and Bitzer, D. L. Advances in computer-based education. Science, 1970, 167, 1582-1590.

Battle, F. L. An analog input interface for the PLATO IV terminal. Urbana, Ill.: Master's thesis, University of Illinois at Urbana-Champaign, Department of Electrical Engineering, 1974.

Bruner, J. Beyond the information given. New York: Norton, 1973.

Carbonell, J. AI in CAI: An artificial intelligence approach to computer-aided instruction. IEEE, Transactions on Man-Machine Systems, 1970, MMS-11, 190-202.

Cochran, W. G. The comparison of percentages in matched samples. Biometrika, 1950, 37, 256-266.

Crowder, N. A. Automatic tutoring by intrinsic programming. In A. A. Lumsdaine and R. Glaser (Eds.); Teaching Machine and Programmed Learning: A Science Book. Washington, D. C.: National Education Association, Department of Audio-Visual Instruction, 1960.

Feurzeig, W. Automated instructional monitors for complex operational tasks. Cambridge, Mass.: Bolt, Beranek and Newman, Report no. 2272, October 1971.

Goldstein, I., Goldstein, G. Flying in your head. Cambridge, Mass.: Massachusetts Institute of Technology, Artificial Intelligence Laboratory, LOGO working paper No. 3, 1972.

Jacobs, R. S. and Roscoe, S. N. Simulator cockpit motion and the transfer of initial flight training. Dallas, Texas: Proceedings of the Nineteenth Annual Meeting of the Human Factors Society. October 1975.

Jacobs, R. S. Simulator cockpit motion and the transfer of initial flight training. Urbana, Ill.: Doctoral dissertation, University of Illinois Aviation Research Laboratory, Technical Report Number ARL-76-8/AFOSR-76-4, June 1976.

Koonce, J. M. Effects of ground based aircraft simulation motion conditions upon prediction of pilot proficiency. Urbana, Ill.: Doctoral dissertation, University of Illinois, Aviation Research Laboratory, Technical Report No. ARL-74-5/AFOSR-74-3, 1974.

Lyman, E. PLATO curriculum materials. Urbana, Ill.: University of Illinois, Computer-based Education Research Laboratory, Report No. X-41 #2, July 1974.

Myers, J. L. Fundamentals of Experimental Design. Boston, Mass.: Allyn and Bacon, Second Edition, 1972.

Neal, J. Electronic laboratory instruction using the CGE-PLATO laboratory station. Urbana, Ill.: University of Illinois. Computer-based Education Research Laboratory, MTC Report 5, December 1974.

Papert, S. A. Teaching children thinking. Programmed Learning and Instructional Technology, 1972, 9, 5.

Papert, S. A. Teaching children to be mathematicians versus teaching about mathematics. International Journal of Mathematical Education and Science Technology, 1972, 3, 249-262.

Papert, S. A. and Solomon, C. Twenty things to do with a computer. Educational Technology, 1972, 12, 4.

Parker, L., and Voss, J. PLATO and driver services. Springfield, Ill.: Department of Motor Vehicles Driver Services Division, 1975.

Povenmire, H.K., and Roscoe, S. N. Incremental transfer effectiveness of a ground-based general aviation trainer. Human Factors, 1973, 15, 6, 534, 542.

Prather, D. Prompted mental practice as a flight simulator. Journal of Applied Psychology, 1973, 57, 3, 353-355.

Quillian, M. R. Semantic memory. In M. Minsky (Ed.), Semantic Processing. Cambridge, Mass.: M. I. T. Press, 1968.

Quillian, M. R. The teachable language comprehender: A simulation program and theory of language. Communications of the ACM, 1969, 12, 459-476.

Roscoe, S. N. Incremental Transfer Effectiveness. Human Factors, 1971, 13, 6, 561-567.

Smallwood, R. D. A decision structure for teaching machines. Cambridge, Mass.: Massachusetts Institute of Technology, Research Monograph, 1962.

Smith, S. G. The use of computers in the teaching of organic chemistry. Journal of Chemical Education, 1970, 47, 9.

Smith, S. G., Ghesquiere, J. R., Avner, R. A. The use of computers in the teaching of chemistry. Journal of Chemical Education, 1974, 51, 6.

SOUPAC Program Description. Urbana, Ill.: University of Illinois Computer Services Office. November 1974.

Trollip, S. R. System performance and student evaluation in a complex computer-based procedural training program. Urbana, Ill.: Master's thesis, University of Illinois Aviation Research Laboratory, Technical Report No. ARL-75-4/AFOSR-75-1, April 1975.

Trollip, S. R. and Ortony, A. Complex skill training in computer-assisted instruction. Instructional Science. In press.

Trollip, S. R. and Roscoe, S. N. Computer-assisted instruction in pilot training and certification. Los Angeles, Cal.: Proceedings of the Sixteenth Annual Meeting of the Human Factors Society. October 1972.

Warnock, E. H., and Collins, A. M. Research on semantic nets. Cambridge, Mass.: Bolt, Beranek, and Newman, Report No. 2555, May 1973.

Winer, B. J. Statistical Principles in Experimental Design. New York: McGraw-Hill, 1975.